



What Makes Mental Associations Personal or Extra-Personal? Conceptual Issues in the Methodological Debate about Implicit Attitude Measures

Bertram Gawronski,* Kurt R. Peters and Etienne P. LeBel
University of Western Ontario

Abstract

Over the last decade, a new class of indirect measurement procedures has become increasingly popular in many areas of psychology. However, these *implicit* measures have also sparked controversies about the nature of the constructs they assess. One controversy has been stimulated by the question of whether some implicit measures (or implicit measures in general) assess extra-personal rather than personal associations. We argue that, despite empirical and methodological advances stimulated by this debate, researchers have not sufficiently addressed the conceptual question of how to define extra-personal in contrast to personal associations. Based on a review of possible definitions, we argue that some definitions render the controversy obsolete, whereas others imply fundamentally different empirical and methodological questions. As an alternative to defining personal and extra-personal associations in an objective sense, we suggest an empirical approach that investigates the meta-cognitive inferences that make a given association subjectively personal or extra-personal for the individual.

Despite the widespread use of self-report measures in various areas of psychology, many researchers tend to be skeptical about their practical usefulness for answering basic psychological questions. For instance, self-report measures have often been criticized for their susceptibility to self-presentation or socially desirable responding (Fazio, Jackson, Dunton, & Williams, 1995). In addition, some mental processes may be unconscious and therefore inaccessible to self-report (Nisbett & Wilson, 1977). For these reasons, psychologists have developed alternative measurement procedures with the goal of eliminating the aforementioned problems.

Over the last decade, a particular class of indirect measurement procedures has become increasingly popular in many areas of psychological research (for reviews, see Fazio & Olson, 2003; Petty, Fazio, & Briñol, forthcoming 2008; Wittenbrink & Schwarz, 2007). In contrast to standard self-report measures, these so-called *implicit measures* are based on experimental

paradigms derived from cognitive psychology, such as sequential priming (Neely, 1977) or response interference tasks (Kornblum, Hasbroucq, & Osman, 1990). The most prominent examples of these measures include Greenwald, McGhee, and Schwartz's (1998) implicit association test (IAT) and Fazio et al.'s (1995) affective priming task. Other examples include semantic priming (Wittenbrink, Judd, & Park, 1997), the go/no-go association task (Nosek & Banaji, 2001), the affect misattribution procedure (Payne, Cheng, Govorun, & Stewart, 2005), and the extrinsic affective Simon task (De Houwer, 2003).

Although implicit attitude measures have been adopted in virtually all sub-disciplines of psychology, and although these measures have proven their usefulness in predicting judgmental biases and overt behavior (Fazio & Olson, 2003), the particular nature of the constructs assessed by these measures is still a subject of debate (Gawronski & Bodenhausen, 2007; Gawronski, LeBel, & Peters, 2007). One controversy has been stimulated by the question of whether implicit measures assess extra-personal rather than personal associations: that is, mental associations that, for some reason or other, may not be considered as part of a person's true self. This argument has been made at various levels of generality, ranging from criticisms of particular kinds of implicit measures (e.g., Olson & Fazio, 2004) to implicit measures in general (e.g., Arkes & Tetlock, 2004).

The main goal of the present article is to scrutinize theoretical conceptualizations of personal and extra-personal associations. Specifically, we argue that the terms *personal* and *extra-personal* have not yet been well defined, which makes it difficult to evaluate empirical and methodological arguments in the ongoing debate about whether implicit attitude measures assess personal or extra-personal associations. More precisely, we claim that the debate regarding personal versus extra-personal associations involves at least three distinct though inter-related components. First, it involves a conceptual component that pertains to how personal versus extra-personal associations are defined in the first place. Second, it involves a methodological component that pertains to procedural aspects of different measurement paradigms, which may determine the relative impact of personal versus extra-personal associations on task performance. Third, it involves an empirical component that pertains to empirically obtained differences between measures that are assumed to assess personal versus extra-personal associations. Drawing on this distinction, we argue that, despite significant advances pertaining to methodological (Olson & Fazio, 2004; Nosek & Hansen, forthcoming 2008) and empirical (De Houwer, Custers, & De Clercq, 2006; Han, Olson, & Fazio, 2006) issues, researchers have not yet sufficiently addressed important aspects pertaining to the first question of how to define personal and extra-personal associations.

To address this limitation, the present article discusses the range and limits of different definitions of personal and extra-personal associations and their respective implications for methodological and empirical aspects

of the current debate. This review will address definitions in terms of (a) endorsement, (b) representation, (c) activation, (d) function, and (e) origin. Although all of these definitions have their own intuitive appeal, we argue that each of them has unique implications for the ongoing debate, with some definitions rendering the controversy obsolete and others implying fundamentally different empirical and methodological questions.

What are Personal and Extra-personal Associations?

Before addressing different definitions of *personal* and *extra-personal* associations, it seems useful to specify the term *association*. For the sake of simplicity, we conceptualize associations as mental links between two concepts in memory. This definition is agnostic about how exactly associations and concepts are represented and therefore does not give any priority to associative network models (e.g., Greenwald et al., 2002; Petty, Briñol, & DeMarree, 2007), localist models (e.g., Read & Miller, 2002; Van Overwalle & Jordens, 2002), or parallel distributed processing models (e.g., Bassili & Brown, 2005; Conrey & Smith, 2007) of mental representation, as they have been discussed in the literature (see also Fazio, 2007). Nevertheless, our specification implies that associations are *mental* entities. For example, the natural co-occurrence between *yellow* and *wasps* has to be mentally represented in order to be assessable by a psychological measure of associations. If this ‘association’ is present only in the natural environment, but not represented in the mind of the participant, it cannot be assessed by any psychological measure (cf. Karpinski & Hilton, 2001).

In addition to specifying the term *association*, it is also important to spell out the criteria on which a satisfying definition of the terms *personal* and *extra-personal* would be based. Drawing on Quine and Ullian’s (1978) principles for evaluating the components of nomological networks, we argue that a useful definition should be sufficiently *precise*, such that it can unambiguously distinguish between objects that do versus do not belong to the category in question. For example, a definition of *animal* may be regarded as insufficiently precise if it cannot determine whether or not dogs would count as animals in terms of the definition. Precision, in turn, is closely related to two intuitive criteria that further contribute to unambiguous classifications. On one hand, a definition should be sufficiently *inclusive*, such that it does not exclude objects that would reasonably be regarded as exemplars of the category in question. For instance, a definition of *animal* may be regarded as insufficiently inclusive if dogs would not count as animals in terms of the definition. On the other hand, a definition should be sufficiently *exclusive*, such that it does not include objects that may be rejected as exemplars of the category in question. For example, a definition of *animal* may be regarded as insufficiently exclusive if tulips would be regarded as animals in terms of the definition. To be sure, the

latter two criteria are vague in that they are based on intuitive preconceptions of the category that needs to be defined. Nevertheless, they often play a significant role when it comes to evaluating the usefulness and plausibility of a given definition. From this perspective, the aforementioned principles are closely related to general meta-theoretical principles, most notably conceptual clarity and logical consistency of the overall nomological network (Quine & Ullian, 1978). These principles also provide the basis for the present endeavor of scrutinizing the range and limits of different definitions of personal versus extra-personal associations.

Endorsement

One possibility for defining personal and extra-personal associations is in terms of endorsement (e.g., Arkes & Tetlock, 2004; Gehring, Karpinski, & Hilton, 2003; Karpinski & Hilton, 2001). According to this definition, a given association is personal when it is explicitly endorsed by an individual. However, an association would be regarded as extra-personal if it is not endorsed. For instance, a mental association between the concepts *African American* and *hostile* would count as a personal association in terms of the definition if it is explicitly endorsed (e.g., if the statement 'African Americans are hostile' is explicitly affirmed). Conversely, it would count as an extra-personal association if it is not endorsed (e.g., if the statement 'African Americans are hostile' is explicitly rejected). This definition is based on a conceptualization of *attitudes* as evaluative judgments (Kruglanski & Stroebe, 2005), implying that only endorsed evaluations represent attitudes in terms of the proposed definition (for alternative conceptualizations of *attitude*, see Banaji, 2001; Eagly & Chaiken, 2007; Fazio, 2007).

A definition in terms of endorsement seems useful in that it provides a precise criterion for regarding a mental association as personal or extra-personal. Empirically, this criterion is reflected in the correspondence between implicit measures and explicit self-reports (for a review, see Hofmann, Gschwendner, Nosek, & Schmitt, 2005), in that implicit measures tapping personal associations should reveal high correspondence to self-reports, whereas implicit measures tapping extra-personal associations may be unrelated to self-reports. However, a definition in terms of endorsement also implies several problems. For instance, Olson and Fazio (2004) explicitly rejected a definition in terms of endorsement, as individuals may regard a given association as accurate, but nevertheless conceal their views in a self-report measure because of self-presentation or social desirability concerns. According to this view, the proposed criterion for personal associations seems overly exclusive, in that some associations would count as extra-personal in terms of the proposed definition even though they may reasonably be regarded as reflecting a person's 'true self.' In addition, a definition in terms of endorsement seems *extrinsic* to the notion of

association, in that the personal or extra-personal character of an association is determined at the level of self-report rather than at the level of the associative representation (Gawronski & Bodenhausen, 2006a). In other words, the proposed definition implies that personal and extra-personal associations may be identical in terms of how they are represented in memory. Rather, their distinct nature would emerge only at the level of explicit judgments, such that some associations are explicitly endorsed whereas others are not.

These considerations have important implications for methodological and empirical questions of the current debate. As implicit measures generally do not assess the explicit *endorsement* of associations (De Houwer, 2006), the methodological question of how procedural variations reflect the relative impact of personal versus extra-personal associations translates into the question of measurement correspondence (i.e., whether procedural variations in implicit measures increase or decrease their correspondence with explicit self-reports). In line with this argument, a series of studies by Olson and Fazio (2004) showed that a 'personalized' variant of the IAT revealed higher correlations to explicit self-report measures compared to the standard variant proposed by Greenwald et al. (1998). Similar conclusions can be drawn for the empirical question regarding observed differences between implicit measures. From the perspective of the present definition, this issue translates into the question of whether effects obtained for self-report measures generalize to implicit measures that have shown high correlations with explicit self-report measures. In line with this reasoning, Han et al. (2006) demonstrated that evaluations that had been endorsed by other individuals in a brief video clip influenced participants' scores on the standard variant of the IAT (Greenwald et al., 1998), although a 'personalized' variant of the IAT (Olson & Fazio, 2004) remained unaffected and in line with participants' self-reported evaluations.

In summary, a definition in terms of endorsement implies that methodological and empirical questions regarding different measurement procedures translate into questions of measurement correspondence (i.e., the correspondence between explicit and implicit measures). However, such a definition seems extrinsic to the notion of association, moving the criterion from an intrinsic feature of mental associations to the extrinsic characteristic of being explicitly endorsed. In other words, whether or not an association is personal or extra-personal is not determined by how associations are represented in memory, but by the endorsement of these associations in self-reported evaluative judgments.

Representation

The extrinsic nature of a definition in terms of endorsement has led some researchers to advocate an intrinsic definition in terms of association representation. For example, with regard to evaluative associations, Olson and

Fazio (2004) argued that extra-personal associations are associations that are available in memory, but which are not part of one's attitude toward a given object. Conversely, personal associations are those that are available in memory and do represent a part of one's attitude. This conceptualization moves the focus onto the definition of *attitude*, which has experienced similar definitional controversies stimulated by the development of implicit attitude measures (Gawronski, 2007). From a general perspective, there are at least three definitions of *attitude* that seem relevant for the present discussion.

First, some researchers define *attitude* as the (endorsed) evaluative judgment of an object (e.g., Kruglanski & Stroebe, 2005). From the perspective of this conceptualization, the proposed definition of personal versus extra-personal associations reverts to the aforementioned one in terms of endorsement (Arkes & Tetlock, 2004; Gehring et al., 2003; Karpinski & Hilton, 2001), thereby transforming the intended intrinsic definition back into an extrinsic one. Needless to say, such a reversal has the same implications that have been discussed for a definition in terms of endorsement.

Second, some researchers define *attitude* as 'a psychological tendency that is expressed by evaluating a particular entity with some degree of favor or disfavor' (Eagly & Chaiken, 1993, p. 1; see also Eagly & Chaiken, 2007). From this perspective, personal associations are defined as those that do result in a tendency to evaluate an object with some degree of favor or disfavor. Conversely, to provide a precise distinction between personal and extra-personal associations, the latter would have to be regarded as those that do not result in such a tendency. The crucial question implied by this conceptualization is whether or not one is willing to interpret an individual's response tendencies on an implicit measure as evaluative tendencies in terms of the proposed definition. If so, the proposed definition of personal associations may be regarded as overly inclusive because any association assessed by an implicit measure would count as a personal association. If not, one would have to specify further criteria that distinguish between response tendencies that do versus do not represent an attitude. To be sure, any equation of undifferentiated measurement scores with a psychological construct may be regarded as overly simplistic, given that performance on implicit measures is typically influenced by multiple distinct processes (Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Payne, 2001). However, to the degree that the relative contribution of these processes can be specified (e.g., Conrey et al., 2005), the lack of process purity does not qualify the argument that a rejection of the proposed equation between response tendencies on implicit measures and evaluative tendencies in terms of Eagly and Chaiken's (1993, 2007) conceptualization requires additional criteria for such a rejection. To our knowledge, such restrictions have not yet been discussed in the literature on attitudes and implicit measures.

Third, some researchers defined *attitudes* as object–evaluation associations (Fazio, 1995, 2007), which is the definition preferred by Olson and Fazio (2004) when they specify extra–personal associations as associations that are available in memory, but which are not part of one’s attitude. However, this definition of personal associations may also be regarded as more inclusive than originally intended. If attitudes are simply defined as object–evaluation associations, then *any* association between an object and a given evaluation would count as an attitude, which would also include ‘extra–personal’ object–evaluation associations. Thus, in order to exclude extra–personal associations as a separate entity, the proposed definition requires further restrictions in the definition of attitudes as object–evaluation associations. Obviously, this restriction cannot by itself refer to the terms *personal* versus *extra–personal*, as such a specification would be circular given the lack of a precise definition of these terms. For instance, Fazio’s (2007) specification of the term *object–evaluation association* as referring to one’s summary evaluation of the object may be regarded as insufficient, given that the term *personal* is simply replaced by the unspecified qualifier *one’s*. Hence, even Fazio’s (2007) specification requires further restrictions to exclude extra–personal associations as a separate entity, as otherwise, any association between an object and a given evaluation would count as an attitude.

Activation

One possible restriction of the definition of attitudes as object–evaluation associations is automatic activation. In line with this restriction, Olson and Fazio (2004) argued that only personal object–evaluation associations tend to be activated automatically, whereas extra–personal object–evaluation associations are not activated automatically.

As with the proposed definition in terms of endorsement, a definition in terms of automatic activation seems useful in that it provides a precise criterion for regarding object–evaluation associations as personal or extra–personal. Nevertheless, even this definition implies some conceptual problems. First, in order to provide an unambiguous distinction between personal and extra–personal associations, the proposed definition would have to classify *any* object–evaluation association as extra–personal if it is not activated automatically. Such a classification of extra–personal associations seems overly inclusive, given that personal object–evaluation associations may vary in terms of their strength and thereby in their capacity to become automatically activated (Fazio, 2007). In other words, the proposed restriction would categorize some object–evaluation associations as extra–personal, although they could reasonably be regarded as reflecting a person’s ‘true’ attitudes or self.

Second, the proposed restriction in terms of automatic activation can acquire multiple meanings when it is applied to the question of whether implicit measures tap personal or extra–personal associations. As outlined

by Bargh (1994), the term *automatic* involves a total of four different components: lack of awareness; lack of intention; lack of controllability; and independence of cognitive resources. Importantly, these characteristics do not necessarily co-vary such that the presence of one property is associated with the presence of the others (see also Moors & De Houwer, 2006). Thus, few, if any, processes meet all four criteria of automaticity (Bargh, 1994; Moors & De Houwer, 2006). From this perspective, it seems useful to limit the proposed restriction to the most relevant features of automaticity instead of including all four characteristics. In line with this argument, Fazio (2007) argued that the most significant criterion is 'inescapability,' which is specified as unintentional activation even when the individual is attempting to engage in some other activity. Based on this specification, personal object–evaluation associations would be defined as those that are activated unintentionally. Conversely, extra-personal object–evaluation associations would have to be regarded as those that require intention to be activated. However, it is important to note that this conclusion stands in direct contrast to a definition in terms of endorsement, which links personal rather than extra-personal associations with intentionality.

Third, even if the notion of automaticity is narrowed to a limited number of the four criteria (e.g., intentionality), whether or not a given measure actually meets these criteria cannot be determined a priori by means of procedural characteristics of the task. As outlined by De Houwer (2006), any claims about features of automaticity represent empirical assumptions, which have to be tested as such. From this perspective, the methodological question of how procedural variations influence the relative impact of personal versus extra-personal associations translates into the question of whether task-related changes in measurement procedures co-vary with construct-related changes in features of automaticity. Moreover, the empirical question regarding observed differences between implicit measures translates into the question of how the obtained features of automaticity vary across measurement procedures, in that a given measure may meet one criterion of automaticity but not others. These differences would then form the premises for understanding empirically observed differences between implicit measures that are assumed to assess personal versus extra-personal associations (e.g., Han et al., 2006). Needless to say, these questions are quite different to the ones implied by a definition in terms of endorsement, which primarily concerns the relation of implicit measures to explicit self-reports. Moreover, given available evidence that many implicit measures can be influenced by intentional memory retrieval (e.g., Blair, Ma, & Lenton, 2001; Gawronski & Bodenhausen, 2005) and other strategic processes (e.g., Fiedler & Blümke, 2005; De Houwer, Beckers, & Moors, 2007; Klauer & Teige-Mocigemba, 2007), these measures would not meet the empirical criterion for assessing personal associations, as defined by the criterion of inescapability.

Function

Another possible restriction on the definition of attitudes as object–evaluation associations pertains to functional properties. In line with such a definition, Olson and Fazio (2004) argued that an important functional value of attitudes resides in their capacity to quickly direct ‘attention, categorization, and ultimately behavior in a manner that maximizes the likelihood of the individual’s experiencing positive outcomes and avoiding negative ones’ (p. 655). From this perspective, object–evaluation associations would count as personal if they fulfill the aforementioned functions by actually influencing these processes. Conversely, to provide a clear and unambiguous distinction between personal and extra–personal associations, the latter would have to be regarded as those that leave these processes unaffected (see also Nosek & Hansen, forthcoming 2008).

At first glance, a functional specification of attitudes, and thus of personal object–evaluation associations, seems useful as it specifies a precise criterion that distinguishes between personal and extra–personal associations. Nevertheless, a definition of personal associations in terms of functional properties could be regarded as overly inclusive, given that object–evaluation associations reflecting cultural norms may also have the potential to influence the aforementioned processes. This notion plays a significant role in theories of attitude–behavior relations, stating that social norms influence the formation of behavioral intentions, and thereby actual behavior, above and beyond attitudes (for a review, see Ajzen & Fishbein, 2005). From this perspective, object–evaluation associations reflecting cultural norms would also count as personal in terms of a functional definition, although such associations might be regarded as not reflecting a person’s ‘true self.’

The proposed functional criterion also seems problematic when it is applied to the empirical and methodological components of the debate. First, as the proposed definition of extra–personal associations implies a general (i.e., unqualified) null effect in the prediction of attention, categorization, and behavior, it would be susceptible to all of the problems attendant to the interpretation of null effects. For example, in the case of implicit attitude measures, null effects in the prediction of behavior could reflect either reliable assessment of extra–personal associations or unreliable assessment of personal associations. Second, a definition in terms of the proposed functions may render the debate obsolete, if there is no a priori criterion that excludes an individual’s responses on the implicit measure as one particular example of behavior. Strictly speaking, implicit measures do not *directly* assess associations. Rather, these measures assess mental constructs *indirectly* by means of behavioral responses that are presumably influenced by associations. From this perspective, a definition in terms of functional properties would make the controversy obsolete, as any association would have to be regarded as personal as soon as it somehow influences behavioral responses on an implicit measure.

In response to these arguments, one could object that a definition in terms of function does not include any type of behavior, as implied by our discussion, but only behavior ‘that maximizes the likelihood of the individual’s experiencing positive outcomes and avoiding negative ones’ (Olson & Fazio, 2004, p. 655). Given this restriction, personal object–evaluation associations could be defined as those that are beneficial for achieving these outcomes, whereas extra-personal associations might be regarded as those that do not fulfill this function. Although this additional restriction resolves most of the aforementioned issues, it still implies a conceptual problem. In general, the valence of a given outcome is not defined a priori, but only with reference to a particular goal. Thus, one and the same behavior can lead to positive or negative outcomes, depending on the particular goal that is used to evaluate the behavior. For instance, taking heroin may be regarded as functional in its capacity to produce positive outcomes, if the goal is to reduce heroin-related drug cravings. However, taking heroin can also be regarded as dysfunctional in its capacity to produce negative outcomes, if the goal is to maintain one’s health. Hence, the proposed restriction to behaviors that are functional in achieving positive outcomes and avoiding negative ones remains insufficiently precise, as the valence of a given outcome, and thus the functionality of the behavior leading to that outcome, is not defined without reference to a particular goal. Moreover, a functional distinction between ‘personal’ and ‘extra-personal’ associations could even be conceived of as a temporal division *within* the person (rather than *between* the person and an extra-personal entity), such that behaviors based on activated associations at one time are not in the interests of the same individual at a later time (Ainslie, 2001).

A possible objection against this argument is that our treatment of positive and negative outcomes as objective properties may be overly restrictive, such that the functional value of attitudes may simply reside in their capacity to ease approach–avoidance decisions (rather than to promote positive and prevent negative outcomes in an objective sense). However, even this limitation seems overly inclusive, as cultural norms or evaluative responses of other individuals may fulfill the same function. This conclusion echoes our earlier argument that social norms tend to influence intentions and behavior above and beyond attitudes (for a review, see Ajzen & Fishbein, 2005), making the two constructs indistinguishable from a mere functional point of view. Given these limitations, a functional restriction of the term *attitude* seems insufficient for a representational definition of personal and extra-personal associations.

Origin

A fifth possible definition of personal versus extra-personal associations refers to the origin of a given association. In line with such a description, some researchers have linked the notion of extra-personal associations to

cultural or normative influences. For instance, in the domain of evaluative associations, extra-personal associations have sometimes been conceptualized as associations that reflect evaluations endorsed by other individuals, which may or may not correspond to one's personal evaluations (e.g., Han et al., 2006). In a similar vein, other researchers have described extra-personal associations as those stemming from empirically observed 'associations' in one's environment, which also may or may not reflect one's personal beliefs (e.g., Arkes & Tetlock, 2004; Gehring et al., 2003; Karpinski & Hilton, 2001).

Although this description may seem intuitively plausible, we consider it the least defensible. From a theoretical perspective, a definition in terms of origins can be translated into one referring to the causes that have created a given association in the first place (e.g., evaluations endorsed by others or observed 'associations' in the environment). If the cause is external to the individual, the association can be regarded as extra-personal. In contrast, if the cause is internal to the individual, the association would count as personal. As outlined by Olson and Fazio (2004), the crucial problem with this definition is that there is no a priori criterion that would specify where exactly one has to stop in the chain of causes and effects (see also Banaji, 2001; Banaji, Nosek, & Greenwald, 2004). Similar to the problem of causal inference in attribution research, any internal (or intra-personal) event that is used to explain a given outcome can be further explained by a preceding external (or extra-personal) event (Malle, 1999). For instance, the 'personal' inference that the 'extra-personal' association between the concepts *African American* and *hostile* is inaccurate could always be related to an external event, such as education in egalitarian values or exposure to friendly African Americans. Frankly, any mental association could be somehow related to an external event, which implies that all associations would have to be regarded as extra-personal according to the proposed definition, making the distinction between personal and extra-personal associations obsolete.

Are There Alternative Approaches?

Our analysis suggests that defining personal and extra-personal associations is not an easy task. Nevertheless, one could object that our discussion of descriptive, categorical definitions is too restrictive, in that it ignores possible alternative approaches to understanding personal and extra-personal associations. In the following sections, we discuss four such approaches: (a) empirical description, (b) operational definitions, (c) combined criteria, and (d) continuous definitions.

Empirical description

In response to the problems associated with the reviewed definitions, one could propose that all of these descriptions can be easily changed into

empirical assumptions about personal and extra-personal associations. For instance, one could argue that personal, in contrast to extra-personal, associations are more likely to be explicitly endorsed; more likely to be activated automatically; more likely to serve the functions of directing attention, categorization, and behavior; and more likely to stem from personal inferences or experiences rather than inferences or experiences by other individuals. From this perspective, the reviewed statements about personal and extra-personal associations would not represent definitions, as suggested in our analysis. Instead, they would represent empirical assumptions about personal and extra-personal associations that can be tested just like any other empirical assumption. In fact, one could argue that all of the reviewed descriptions of personal and extra-personal associations may have never been intended as definitions but as empirical statements. In this case, most of the problems highlighted by our analysis would become irrelevant, as they primarily apply to definitions but not to empirical statements. At the same time, however, the deplored lack of a definition would become even more problematic, as it would imply that two undefined labels have been used to interpret empirical results obtained with implicit measures. In other words, an empirical reinterpretation of the reviewed descriptions still requires a clear specification of what personal and extra-personal associations are to begin with. Without such a definition, it would remain unclear what entities these statements are referring to, making empirical tests of these assumptions strictly impossible. Hence, it seems important to clearly define the notions of personal and extra-personal associations and to indicate whether a given statement about personal or extra-personal associations should be regarded as part of their definition or as an empirical assumption. If a statement is intended as part of a definition, it may also be useful to take into account the discussed criteria (i.e., precision, inclusiveness, and exclusiveness) for a useful definition.

Operational definitions

Another possibility to deal with the problems implied by the reviewed *descriptive* definitions is to adopt an *operational* definition of personal and extra-personal associations. From this perspective, the two types of associations are not defined by descriptive statements regarding their nature, but by the measurement procedures that are used to assess personal and extra-personal associations. For instance, personal associations could be defined as the associations that are assessed with particular measurement procedures, such as the personalized variant of the IAT (Olson & Fazio, 2004) or affective priming (Fazio et al., 1995). Conversely, extra-personal associations could be defined as those assessed by other types of measures, such as the cultural norms variant of the IAT (Yoshida, Peach, Spencer, & Zanna, 2007). As with the aforementioned empirical reinterpretation, an operational approach to defining personal and extra-personal associations

would circumvent most of the problems outlined for the reviewed descriptive definitions. Nevertheless, even an operational definition of the two types of associations is not without problems.

First, any operational definition implies the risk of infinitely multiplying the number of constructs, which would violate Occam's principle of parsimony, stating that the number of postulated constructs should be kept as low as possible (see Quine, 1963). This issue has its precedent in the literature on implicit measures, in that explicit and implicit measures have sometimes been postulated to assess two distinct attitudes that are stored independently in memory (e.g., Wilson, Lindsey, & Schooler, 2000). This assumption has been criticized by several scholars, who argued that dissociations between the two types of measures do not necessarily indicate the existence of two distinct attitudes in memory (e.g., Fazio & Olson, 2003; Gawronski & Bodenhausen, 2006b; Petty et al., 2007). Instead, these dissociations can be easily explained by process models of attitudes and evaluation, which propose only a single attitude to explain the same set of empirical findings.

Second, any equation of a construct with the outcome of a measurement procedure necessarily treats the measurement procedure as process pure. However, this assumption of process purity stands in contrast to the insight that virtually any measurement procedure involves the operation of multiple distinct processes. This conclusion also applies to implicit measures, which have been shown to be affected by several processes other than activated associations (e.g., Conrey et al., 2005; Payne, 2001). From this perspective, an operational definition of personal and extra-personal associations runs the risk of oversimplification by equating the construct that needs to be defined with a conglomerate of multiple processes that goes far beyond the to-be-defined construct.

Combined criteria

Another possible objection against our theoretical analysis is that even though none of the discussed criteria seems sufficient for providing a satisfying distinction between personal and extra-personal associations, a combination of multiple necessary criteria could possibly meet this requirement. This objection echoes our earlier argument regarding a definition in terms of representation, such that the implications of this definition depend on the preferred conceptualization of *attitude*. Moreover, given a conceptualization of attitudes as object-evaluation associations (Fazio, 1995, 2007), a definition in terms of representation still needs to specify further restrictions for the term object-evaluation association, as otherwise any object-evaluation association would count as personal in terms of the proposed definition. Two possible restrictions that have been discussed in this context are activation and function. Thus, the objection that our discussion treats the five criteria as mutually exclusive is not entirely

correct, given the discussed necessity of combining criteria for a definition in terms of representation.

Notwithstanding the importance of further restrictions for a definition in terms of representation, it is important to note that one must be cautious in combining criteria, as such combination can result in empty categories. For instance, definitions that imply that all associations are personal (e.g., function) or extra-personal (e.g., origin) do not leave any room for the respective alternative category. Thus, a combination of such contradictory definitions would result in empty categories for both personal and extra-personal associations, such that no association can be regarded as either personal or extra-personal. Needless to say, such a specification is even more exclusive than the individual definitions that underlie the combination of criteria.

Continuous definitions

Another possible objection is that our emphasis on distinct categories may be too rigid, given that most conceptual distinctions in social and cognitive psychology may actually reflect continua. For instance, the notion of cognitive elaboration is obviously a matter of degree rather than two distinct categories of high versus low elaboration. In a similar vein, the distinction between personal and extra-personal associations may be a matter of degree, such that associations may be more or less personal. Notwithstanding the intuitive plausibility of this objection, it seems important to distinguish between conceptual and empirical aspects of a definition. To be sure, a theoretically specified distinction between two concepts may be a matter of degree at the empirical level, such that the presence of their defining features can be more or less pronounced. However, this does not imply that the defining features themselves are a matter of degree at the conceptual level. For instance, although cognitive elaboration may be a matter of degree at the empirical level, the term *cognitive elaboration* still requires a precise and unambiguous definition of what cognitive elaboration is in the first place. Thus, although theoretical distinctions between two constructs may turn out to be continuous at the empirical level, their definitions nevertheless have to provide a clear distinction between the two at the conceptual level. This is also true for any definition of personal and extra-personal associations, which should provide an unambiguous, categorical distinction at the conceptual level, although this categorical distinction may manifest itself in a continuous manner at the empirical level.

The Moral Side of the Debate

Irrespective of our particular conclusions, it is important to note that definitional issues cannot be resolved empirically. Simply put, the question of how to define psychological concepts is one of consensus, not empirical

discovery, and in this sense must be resolved through dialogue. Needless to say, the primary goal of such dialogues should be conceptual clarity and logical consistency of the overall nomological network. This has been the focus adopted in the present review, in which we tried to illustrate the range and limits of different possible definitions from a conceptual point of view. However, with regard to the debate about personal and extra-personal associations, the scientific dialogue is likely influenced by a second, non-scientific concern. Given that any description of a given association as personal or extra-personal involves a notion of personal responsibility, the scientific debate may be shaped by moral considerations that go far beyond the quest for conceptual clarity and logical consistency. Such moral considerations have become increasingly important with evidence of the impact of associative knowledge on overt behavior, even when these behaviors may be disregarded as driven by 'extra-personal' forces (Bargh & Ferguson, 2000; Strack & Deutsch, 2004). Given these findings, the question of how to define personal and extra-personal associations involves a strong moral component, in that any answer to this question has the potential to influence judgments of personal responsibility. A striking example of this issue is the current debate of whether police officers should be held responsible for the stronger tendency to pull the trigger of their guns in response to Black compared with White individuals (e.g., Correll, Park, Judd, & Wittenbrink, 2002; Greenwald, Oakes, & Hoffman, 2003; Payne, 2001). Whereas a definition in terms of endorsement may refer to the potential lack of intention in classifying the mental associations leading to this tendency as extra-personal (cf. Arkes & Tetlock, 2004), a definition in terms of activation may cite the same lack of intention as evidence for the personal nature of these associations and thus for the true nature of the person pulling the trigger (cf. Payne, Jacoby, & Lambert, 2005). Needless to say, this discussion goes far beyond the scientific discourse on how to define a psychological construct.

Conclusions

In their discussion of different kinds of implicit attitude measures, Olson and Fazio (2004) noted 'some conceptual fuzziness' (p. 659) pertaining to the distinction between personal and extra-personal associations. The main goal of the present article was to address this fuzziness by scrutinizing different possible definitions of personal versus extra-personal associations. Specifically, we argued that the controversy regarding the nature of the associations assessed by implicit measures involves three distinct though interrelated components: (a) a conceptual component that pertains to the definition of personal versus extra-personal associations; (b) a methodological component that pertains to how procedural variations influence the relative impact of personal versus extra-personal associations on task performance; and (c) an empirical component that pertains to empirically

obtained differences between measures that are assumed to assess personal versus extra-personal associations. Our central argument is that the current debate has primarily focused on the second and third components, but has not sufficiently addressed important questions pertaining to the first one.

To address this concern, we reviewed five possible definitions of personal and extra-personal associations in terms of (a) endorsement, (b) representation, (c) activation, (d) function, and (e) origin. Our discussion revealed that, even though all of these definitions have their own intuitive appeal, each of them has unique implications for the ongoing debate. Whereas some definitions render the controversy obsolete by implying that all implicit measures, by definition, exclusively assess either personal associations (e.g., a definition in terms of function) or extra-personal associations (e.g., a definition in terms of origin), others move the resulting methodological and empirical questions in fundamentally different directions (i.e., correspondence between explicit and implicit measures; features of automaticity). Moreover, the only two definitions that do not seem to make the distinction entirely obsolete (endorsement and activation) are partially inconsistent with each other, in that one of them links personal associations to intentionality (endorsement), whereas the other links personal associations to unintentionality (activation). However, even these two definitions involve significant conceptual problems. Whereas a definition in terms of endorsement employs a criterion that is extrinsic to the notion of mental associations, a definition in terms of activation may be regarded as overly exclusive, given that even personal object-evaluation associations may vary in terms of their strength, and thereby in their potential of becoming automatically activated. In fact, the proposed equation of personal associations with those that are activated automatically seems rather distant from prior discussions in social and cognitive psychology, in which association activation has traditionally been linked to the concepts of availability, accessibility, salience, and applicability (Higgins, 1996; Förster & Liberman, 2007) rather than the personal or extra-personal nature of associations.

What implications do these conclusions have for theoretical and empirical investigations using implicit measures? In our own research, we have rejected the *objective* notion implied by the reviewed definitions of personal and extra-personal associations (e.g., Gawronski & Bodenhausen, 2006a). Instead, we suggest a process-based approach in which the 'personal' or 'extra-personal' character of an association is determined *subjectively* by the individual. Similar to recent developments in the debate on free will, shifting from the question of whether free will exists in an objective sense to the empirical question of how people come to experience free will (Wegner & Wheatley, 1999), the question in our model is not whether personal or extra-personal associations exist in an objective sense, but how people come to consider an association as part of their true self and thus as either personal or extra-personal. From this perspective, the question is

not how to conceptualize personal and extra-personal associations as objective entities of empirical observation, but which processes contribute to people's subjective conclusion that a given association is personal or extra-personal. Although this approach may appear superficially similar to a definition in terms of endorsement, the two approaches differ in that the latter aims at identifying features of associations that make them personal or extra-personal in an objective sense, whereas our approach focuses on the psychological processes that make a given association subjectively personal or extra-personal for the individual.

In our view, a given individual may regard a mental association as personal to the degree that he or she perceives this association as accurate or valid (Gawronski & Bodenhausen, 2006a). Subjective validity, in turn, is assumed to be the result of a process of propositional reasoning, in which the individual assesses the consistency of this association with all other relevant information that is momentarily activated (e.g., Gawronski, Peters, Brochu, & Strack, forthcoming 2008). If the association is consistent with that information, it will likely be regarded as valid. However, if the association is inconsistent with other accessible information, inconsistency has to be resolved to avoid uncomfortable feelings of cognitive dissonance (see Festinger, 1957). Depending on the momentarily employed inconsistency resolution strategy, dissonance may be reduced by a rejection of the association as invalid (Gawronski & Strack, 2004). In such cases, the association itself may be attributed to external sources that could be deemed questionable, which makes that association subjectively 'extra-personal' for the individual. These processes resemble the ones that have recently been investigated under the label *authorship attribution* (e.g., Dijksterhuis, Preston, Wegner, & Aarts, forthcoming 2008). According to this view, the personal or extra-personal character of a given association is not defined by objective features of its representation, as implied by the reviewed attempts to define personal and extra-personal associations in an objective manner. Instead, the personal or extra-personal character of an association is purely subjective, such that it depends on the outcome of meta-cognitive processes of source attribution (i.e., me vs. not me). From this perspective, any use of the terms *personal* or *extra-personal* as objective descriptions of mental entities seems misleading, as it implicitly suggests that there are two types of associations that are inherently distinct.

Despite our rejection of the distinction between personal and extra-personal associations in objective terms, it is important to note that our analysis does not negate meaningful differences between different measurement procedures currently used by psychologists. In fact, there is accumulating evidence that minor variations in measurement procedures can have significant influences on task performance and the resulting scores (e.g., De Houwer et al., 2006; Han et al., 2006; Nosek & Hansen, forthcoming 2008; Olson & Fazio, 2004). However, given the lack of a clear definition, any explanation of these differences as reflecting a differential impact of

personal versus extra-personal associations seems circular. Drawing on proposals emphasizing the role of procedural characteristics of implicit measures (e.g., De Houwer, 2006), we argue that it may be more useful to relate empirically obtained differences between measures to operational features of their measurement procedures. For example, in our own research, we have investigated how the presence versus absence of a response interference component determines the relative impact of accessible content versus accessibility experiences (Gawronski & Bodenhausen, 2005) and enhanced stimulation of associations (Deutsch & Gawronski, forthcoming) on implicit measures. This approach is in line with the methodological and empirical implications of the discussed definition in terms of activation, in that both call for systematic investigations of the relation between operational characteristics of measurement procedures and different features of automaticity. However, the proposed approach differs from a definition in terms of activation, such that it rejects the replacement of these features (e.g., intentionality) with the yet-to-be-defined terms *personal* and *extra-personal*.

Another possibility is implied by our re-interpretation of the terms *personal* and *extra-personal* as referring to the subjective outcome of meta-cognitive inferences rather than objective features of associations. Specifically, one could argue that the outcome of these inferences may be stored in associative memory. Depending on the particular conclusion drawn by an individual (i.e., that a given association is personal or extra-personal), these meta-cognitive associations may then differentially affect measures that have been argued to assess personal versus extra-personal associations (cf. Han et al., 2006; Olson & Fazio, 2004). It is important to note, however, that such meta-cognitive associations may be better described as a person's attitude self-concept (i.e., a stored representation of an individual's meta-cognitive conclusion about the nature of his or her attitude toward a given object). Such attitude self-concepts are distinct from attitudes per se, in that attitude self-concepts are simply descriptive and do not involve an evaluative response to the attitude object. Thus, measures that have been claimed to provide purer reflections of 'personal' associations (e.g., Olson & Fazio, 2004) may in fact assess people's self-concept of their attitudes rather than attitudes per se. To be sure, this conceptual re-interpretation is speculative at this point and future research is needed to determine its validity. In any event, we believe that empirical investigations focusing on the processes underlying implicit measures will provide deeper insights into the mechanisms that are responsible for previously obtained differences between measures relative to labelling these differences as reflecting personal versus extra-personal associations. In addition, research may shift from current attempts to conceptualize personal and extra-personal associations in an objective sense to investigating the meta-cognitive processes that lead people to conclude that a given association is personal or extra-personal.

Acknowledgements

Preparation of this article was supported by grants from the Canada Research Chairs Program (CRC), the Social Sciences and Humanities Research Council of Canada (SSHRC), and the Academic Development Fund of the University of Western Ontario (ADF). We would like to thank Russ Fazio, Mike Olson, and Jim Olson for helpful comments on earlier versions of this article.

Short Biographies

Bertram Gawronski's research lies at the intersection of social and cognitive psychology, addressing the interplay of automatic and controlled processes in social behavior and information processing. He has authored articles in these areas for various top-tier journals, including *Psychological Bulletin*, the *Journal of Personality and Social Psychology*, the *Journal of Experimental Social Psychology*, and the *Personality and Social Psychology Bulletin*, and recently served as a guest editor for a special issue of *Social Cognition* titled 'What is an attitude?'. His current research addresses the interplay of associative and propositional processes in attitude change, which is guided by his associative-propositional evaluation (APE) model developed in collaboration with Galen V. Bodenhausen. His research has been distinguished with several awards, including the 2006 Theoretical Innovation Prize from the Society for Personality and Social Psychology and the 2006 Early Career Award from the International Social Cognition Network. He received his PhD in psychology from Humboldt-University Berlin (Germany), and currently holds a Canada Research Chair in Social Psychology at the Department of Psychology at the University of Western Ontario (Canada).

Kurt Peters is a PhD student in the Department of Psychology at the University of Western Ontario. His research lies at the intersection of social, cognitive, and evolutionary psychology, addressing associative and propositional components of self-representation. He received his BA in computer science and music from Dartmouth College (USA) and his MSc in psychology from the University of Western Ontario (Canada).

Etienne LeBel is a PhD student in the Department of Psychology at the University of Western Ontario. His research lies at the intersection of social and cognitive psychology, addressing the dynamics of self-evaluation, including origins, development, and experiences of self-esteem. He received his BA in psychology from the University of Waterloo (Canada) and his MSc in psychology from the University of Western Ontario (Canada).

Endnote

* Correspondence address: Department of Psychology, University of Western Ontario, Social Science Centre, London, Ontario, Canada N6A 5C2. Email: bgawrons@uwo.ca

References

- Ainslie, G. (2001). *Breakdown of Will*. New York: Cambridge University Press.
- Ajzen, I., & Fishbein, M. (2005). The influence of attitudes on behavior. In D. Albarracín, B. T. Johnson & M. P. Zanna (Eds.), *The Handbook of Attitudes* (pp. 173–221). Mahwah, NJ: Erlbaum.
- 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 III, 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., Nosek, B. A., & Greenwald, A. G. (2004). No place for nostalgia in science: A response to Arkes & Tetlock. *Psychological Inquiry*, **15**, 279–289.
- Bargh, J. A. (1994). The four horsemen of automaticity: Awareness, intention, efficiency, and control in social cognition. In R. S. Wyer & T. K. Srull (Eds.), *Handbook of Social Cognition* (pp. 1–40). Hillsdale, NJ: Erlbaum.
- Bargh, J. A., & Ferguson, M. J. (2000). Beyond behaviorism: On the automaticity of higher mental processes. *Psychological Bulletin*, **126**, 925–945.
- Bassili, J. N., & Brown, R. D. (2005). Implicit and explicit attitudes: Research, challenges, and theory. In D. Albarracín, B. T. Johnson & M. P. Zanna (Eds.), *The Handbook of Attitudes* (pp. 543–567). Mahwah, NJ: Lawrence Erlbaum Associates.
- 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.
- Conrey, F. R., & Smith, E. R. (2007). Attitude representation: Attitudes as patterns in a distributed, connectionist representational system. *Social Cognition*, **25**, 718–735.
- 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.
- Correll, J., Park, B., Judd, C. M., & Wittenbrink, B. (2002). The police officer’s dilemma: Using ethnicity to disambiguate threatening individuals. *Journal of Personality and Social Psychology*, **83**, 1314–1329.
- De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology*, **50**, 77–85.
- De Houwer, J. (2006). What are implicit measures and why are we using them? In R. W. Wiers & A. W. Stacy (Eds.), *The Handbook of Implicit Cognition and Addiction* (pp. 11–28). Thousand Oaks, CA: Sage Publishers.
- De Houwer, J., Beckers, T., & Moors, A. (2007). Novel attitudes can be faked on the Implicit Association Test. *Journal of Experimental Social Psychology*, **43**, 972–978.
- De Houwer, J., Custers, R., & De Clercq, A. (2006). Do smokers have a negative implicit attitude toward smoking? *Cognition and Emotion*, **20**, 1274–1284.
- Deutsch, R., & Gawronski, B. (forthcoming). *When the Method Makes a Difference: Antagonistic Effects on ‘Automatic Evaluations’ as a Function of Task Characteristics of the Measure*. Manuscript submitted for publication.
- Dijksterhuis, A., Preston, J., Wegner, D. M., & Aarts, H. (forthcoming 2008). Effects of subliminal priming of self and God on self-attribution of authorship for events. *Journal of Experimental Social Psychology*.
- Eagly, A. H., & Chaiken, S. (1993). *The Psychology of Attitudes*. Fort Worth, TX: Harcourt Brace Jovanovich.
- Eagly, A. H., & Chaiken, S. (2007). The advantages of an inclusive definition of attitude. *Social Cognition*, **25**, 582–602.
- Fazio, R. H. (1995). Attitudes as object-evaluation associations: Determinants, consequences, and correlates of attitude accessibility. In R. E. Petty & J. A. Krosnick (Eds.), *Attitude Strength* (pp. 247–282). Mahwah, NJ: Erlbaum.
- Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition*, **25**, 603–637.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology*, **54**, 297–327.

- Fazio, R. H., 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*. Evanston, IL: Row Peterson.
- Fiedler, K., & Blümke, M. (2005). Faking the IAT: Aided and unaided response control on the Implicit Association Tests. *Basic and Applied Social Psychology*, **27**, 307–316.
- Förster, J., & Liberman, N. (2007). Knowledge activation. In A. W. Kruglanski, & E. T. Higgins (Eds.), *Social Psychology: Handbook of Basic Principles* (2nd edn, pp. 201–231). New York: Guilford.
- Gawronski, B. (2007). Attitudes can be measured! But what is an attitude? *Social Cognition*, **25**, 573–581.
- Gawronski, B., & Bodenhausen, G. V. (2005). Accessibility effects on implicit social cognition: The role of knowledge activation versus retrieval experiences. *Journal of Personality and Social Psychology*, **89**, 672–685.
- Gawronski, B., & Bodenhausen, G. V. (2006a). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin*, **132**, 692–731.
- Gawronski, B., & Bodenhausen, G. V. (2006b). Associative and propositional processes in evaluation: Conceptual, empirical, and meta-theoretical issues. Reply to Albarracín, Hart, and McCulloch (2006), Kruglanski and Dechesne (2006), and Petty and Briñol (2006). *Psychological Bulletin*, **132**, 745–750.
- Gawronski, B., & Bodenhausen, G. V. (2007). What do we know about implicit attitude measures and what do we have to learn? In B. Wittenbrink & N. Schwarz (Eds.), *Implicit Measures of Attitudes* (pp. 265–286). New York: Guilford Press.
- Gawronski, B., LeBel, E. P., & Peters, K. R. (2007). What do implicit measures tell us? Scrutinizing the validity of three common assumptions. *Perspectives on Psychological Science*, **2**, 181–193.
- Gawronski, B., Peters, K. R., Brochu, P. M., & Strack, F. (forthcoming 2008). Understanding the relations between different forms of racial prejudice: A cognitive consistency perspective. *Personality and Social Psychology Bulletin*.
- Gawronski, B., & Strack, F. (2004). On the propositional nature of cognitive consistency: Dissonance changes explicit, but not implicit attitudes. *Journal of Experimental Social Psychology*, **40**, 535–542.
- Gehring, W. J., Karpinski, A., & Hilton, J. L. (2003). Thinking about interracial interactions. *Nature Neuroscience*, **6**, 1241–1243.
- Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, **109**, 3–25.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. K. L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, **74**, 1464–1480.
- Greenwald, A. G., Oakes, M. A., & Hoffman, H. (2003). Targets of discrimination: The effect of race on detecting presence of weapons. *Journal of Experimental Social Psychology*, **39**, 399–405.
- Han, H. A., Olson, M. A., & Fazio, R. H. (2006). The influence of experimentally-created extrapersonal associations on the Implicit Association Test. *Journal of Experimental Social Psychology*, **42**, 259–272.
- Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social Psychology: Handbook of Basic Principles* (pp. 133–168). New York: Guilford.
- Hofmann, W., Gschwendner, T., Nosek, B. A., & Schmitt, M. (2005). What moderates implicit-explicit consistency? *European Review of Social Psychology*, **16**, 335–390.
- Karpinski, A., & Hilton, J. L. (2001). Attitudes and the Implicit Association Test. *Journal of Personality and Social Psychology*, **81**, 774–788.
- Klauer, K. C., & Teige-Mocigemba, S. (2007). Controllability and resource dependence in automatic evaluation. *Journal of Experimental Social Psychology*, **43**, 648–655.
- Kornblum, S., Hasbroucq, T., & Osman, A. (1990). Dimensional overlap: Cognitive basis for stimulus-response compatibility – a model and taxonomy. *Psychological Review*, **97**, 253–270.

- Kruglanski, A. W., & Stroebe, W. (2005). The influence of beliefs and goals on attitudes: Issues of structure, function, and dynamics. In D. Albarracín, B. T. Johnson & M. P. Zanna (Eds.), *The Handbook of Attitudes* (pp. 323–368). Mahwah, NJ: Lawrence Erlbaum Associates.
- Malle, B. (1999). How people explain behavior: A new theoretical framework. *Personality and Social Psychology Review*, **3**, 23–48.
- Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, **132**, 297–326.
- Neely, J. H. (1977). Semantic priming and retrieval from lexical memory: Roles of inhibition-less spreading activation and limited-capacity attention. *Journal of Experimental Psychology: General*, **106**, 226–254.
- Nisbett, R. E., & Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review*, **84**, 231–259.
- Nosek, B. A., & Banaji, M. R. (2001). The go/no-go association task. *Social Cognition*, **19**, 625–666.
- Nosek, B. A., & Hansen, J. (forthcoming 2008). The associations in our head belong to us: Searching for attitudes and knowledge in implicit evaluation. *Cognition and Emotion*.
- Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extra-personal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology*, **86**, 653–667.
- 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., Cheng, S. 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.
- 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). New York: Oxford University Press.
- 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**, 657–686.
- Petty, R. E., Fazio, R. H., & Briñol, P. (Eds.) (forthcoming 2008). *Attitudes: Insights from the New Implicit Measures*. Mahwah, NJ: Erlbaum.
- Quine, W. V. O. (1963). On simple theories in a complex world. *Synthese*, **15**, 103–106.
- Quine, W. V. O., & Ullian, J. S. (1978). *The Web of Belief* (2nd edn). New York: McGraw-Hill.
- Read, S. J., & Miller, L. C. (2002). Virtual personalities: A neural network model of personality. *Personality and Social Psychology Review*, **6**, 357–369.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, **8**, 220–247.
- Van Overwalle, F., & Jordens, K. (2002). An adaptive connectionist model of cognitive dissonance. *Personality and Social Psychology Review*, **6**, 204–231.
- Wegner, D. M., & Wheatley, T. (1999). Apparent mental causation: Sources of the experience of free will. *American Psychologist*, **54**, 480–492.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, **107**, 101–126.
- Wittenbrink, B., & Schwarz, N. (Eds.) (2007). *Implicit Measures of Attitudes: Procedures and Controversies*. New York: Guilford Press.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationships with questionnaire measures. *Journal of Personality and Social Psychology*, **72**, 262–274.
- Yoshida, E., Peach, J. M., Spencer, S. J., & Zanna, M. P. (2007). *How Social Representations Become Automatic: The Measurement and Impact of Implicit Cultural Norms*. Unpublished manuscript, University of Waterloo.