Who Starts the Wave? Let's not forget the role of the individual

Javier A. Granados Samayoa Russell H. Fazio

The Ohio State University

Author Note

Correspondence concerning this article should be addressed to Russell H. Fazio, Department of Psychology, The Ohio State University, 1835 Neil Avenue, Columbus, Ohio 43210. Email: fazio.11@osu.edu.

Implicitly-measured intergroup bias is a popular and important topic of study. As a result, the field has witnessed the accumulation of an impressive number of investigations, some of which have yielded seemingly puzzling findings. In the target article, Payne, Vuletich, and Lundberg (this issue) present a theoretical framework that aims to make sense of these puzzles by suggesting that implicit measures of intergroup bias are reflections of situational contexts, not the evaluations of individuals. That is, instead of conceptualizing implicitly-measured intergroup bias as accessible attitudes toward a social group that reside within the individual, Payne et al. view it as reflecting the accessibility of group-evaluation associations that are determined by the situation in which the individual is immersed.

Highlighting Two Points of Agreement

Before jumping to the more substantive commentary, we would like to first highlight a couple of issues that have arisen in the literature and that were discussed in the target article. We believe these points of agreement are sufficiently important that they merit reiterating here. First, implicit bias is sometimes defined as an attitude about a social category that is unconscious in the sense that it cannot be accessed through introspection (i.e., bias that people do not know they have; Kang et al., 2012). The dissociation sometimes observed between implicit and explicit measures of attitudes is cited as evidence for this

view (Greenwald & Nosek, 2008). While this perspective remains popular (e.g., Capers, Clinchot, McDougle, & Greenwald, 2017), its appropriateness has been questioned. As Fazio and Olson (2003) have argued, nothing about the measures employed in this line of research guarantees that individuals lack awareness of the associations being assessed. For example, an individual who has more difficulty pairing a given social group with the pleasant category than with the unpleasant category is not necessarily unaware of their relatively negative associations with that group. In fact, as Payne and colleagues highlight, when individuals were asked about their "gut reactions," these reports and implicitly-measured bias scores were significantly correlated (Ranganath, Smith, & Nosek, 2008). This finding suggests that individuals can consciously access the associations that implicit measures of bias assess, provided that the right question is asked (see also Hahn, Judd, Hirsh, & Blair, 2014).

Thus, the use of the term "implicit" when referring to the attitudes themselves does not seem to be an appropriate descriptor of what is captured by measures like the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), the Affect Misattribution Procedure (AMP; Payne, Cheng, Govorun, & Stewart, 2005), or the Evaluative Priming Measure (EPM; Fazio, Jackson, Dunton, & Williams, 1995). The automatic activation of attitudes should not be equated with individuals' lack of awareness of the attitude. Following Fazio and Olson (2003), we advocate referring to the measure and not

the construct as implicit. Hence, we prefer the use of the term implicitly-measured attitudes when discussing what is assessed by the IAT, AMP, or EPM.[1] Individuals may lack awareness that their attitudes are being measured because they are not directly asked to report them, but there is a dearth of evidence that they lack awareness of holding such attitudes. Moreover, we prefer to view "implicit bias" as an effect that individuals' attitudes may have on their judgments or behavior in a given situation without their awareness of their having been influenced by their automatically-activated attitudes.

The second issue that we would like to highlight involves the source of the associations captured by implicit measures. Importantly, the argument assumes that implicitly-measured intergroup bias can reside within an individual in the same way other attitudes do (a viewpoint we will expand upon later). As Payne and colleagues point out, some theorists have postulated that early-life experiences are a major source of the associations captured by implicit measures of attitudes (e.g., Rudman, 2004). Contrary to this perspective, however, one study found that recent religious experiences and behaviors predicted implicitly-measured attitudes toward religion, while early-life experiences and behaviors did not (Castelli, Carraro, Gawronski, & Rava, 2010). More relevant to the issue of intergroup bias, implicitly-measured attitudes toward Black individuals in a

[1] When referring to instances in which the attitude object is a social category, we use the term implicitly-measured intergroup bias.

sample of college freshmen were significantly related to the favorability of their interactions with Black individuals during high school, but not middle school or elementary school (Towles-Schwen & Fazio, 2001). Thus, the data are consistent with the idea that the associations indexed by implicit measures need not stem from early childhood, but instead can be more strongly influenced by recent experiences.

Abstract Categories?

The major thesis offered by Payne et al. is that implicit measures of intergroup bias reflect situations more than persons. Noting that attitudes can exist as chronically accessible linkages to objects or people and in that way "reflect stable aspects of the person" (p. 17), they are careful to restrict their proposition to the implicit measurement of intergroup biases. They argue that a social category is a more abstract attitude object than, say, a given political candidate, rendering it more dependent on context. Given that differentiating the situationally-driven form of implicitly-measured intergroup bias from other dispositionally-driven phenomena is a crucial question for the theory, we were left wanting a deeper exploration of this issue. Is the same reasoning applicable to all categories of people? Does it apply equally to social categories that are commonly regarded as potentially tainted by inappropriate stereotyping and prejudice (e.g., African-Americans, women, Muslims) as to ones for which social desirability concerns

are less likely to arise (e.g., Democrats, Republicans, lawyers, engineers)? Are the latter any less abstract? If not, then what additional attributes characterize intergroup biases that are likely to be tied more to situations than persons? Similarly, would the same reasoning apply to abstract categories of objects (e.g., guns) or animals (e.g., cats) that themselves include a diverse array of exemplars? We suspect not. Whether our intuitions are correct or not, these questions highlight the need for a more detailed theoretical framework regarding the domains likely to be characterized by strong situational influences.

Our Perspective on Implicitly-Measured Intergroup Bias

That aside, the authors lay out a very thought-provoking framework. Specifically, we consider the idea that what has previously been thought of as error variance when measuring intergroup biases implicitly could represent systematic situational variance to be a very valuable insight. However, we do not view situations as the only source of systematic variance captured by implicit measures of intergroup bias. The main thesis of our commentary is that while the situation may be a powerful force, implicit measures of intergroup bias do reflect personal attitudes, at least for some individuals and some measures. Moreover, the impact of these personal attitudes may be attenuated by such factors as the motivation and opportunity to control automatically-activated attitudes – factors that can moderate the attitude-to-behavior process. In the paragraphs that follow,

we will argue that both the situation and the person play important roles in understanding implicitly-measured intergroup bias and its consequences.

Who Starts the Wave?

In the target article, the authors illustrate their view of implicit bias by considering the behavior of fans when "the wave" has broken out in a stadium. They argue that the most accurate prediction of behavior is provided by measuring whether the wave is going through the stadium at the time, rather than measuring attributes of the individuals in the stadium. But what if our interest is not in predicting whether or when a given fan will stand up upon the arrival of the wave, but rather in predicting who will start the wave? Or, what if our interest is in who might actually resist the wave? Here, attributes of the individual would be the best predictors. Thus, it is important to not forget the critical role that can be occupied by the individual.

Considerable evidence exists regarding the predictive validity of the estimates of individuals' attitudes derived from implicit measures. For example, implicit measures of attitudes toward one's romantic partner have been found to prospectively predict long-term changes in marital satisfaction (McNulty, Olson, Meltzer, & Shaffer, 2013) as well as dissolution of the relationship (Lee, Rogge, & Reis, 2010). Similar findings have been observed with respect to an evaluative priming measure of automatically-activated racial attitudes and interracial

dormitory roommate relationships (Towles-Schwen & Fazio, 2006). White freshmen who had been randomly assigned to share a room with Black freshmen were brought into lab early in the Fall semester to have their attitudes toward Black individuals assessed via the EPM. At the end of the academic year, the researchers ascertained the longevity of the roommate relationship via university records. The results revealed that the implicit measure significantly predicted the duration of the roommate relationship. The interracial roommate relationships were more likely to dissolve, and to do so earlier, when the White roommate was characterized by more negative racial attitudes.

What is particularly noteworthy about these results is that the implicit measures predicted temporally distant outcomes, which suggests that what the implicit measure assesses is a chronically-accessible evaluation (i.e., a personal attitude). Moreover, in the context of the regular interactions involved in long-term interpersonal relationships, individuals are unlikely to have the consistent opportunity to control the influence of their automatically-activated attitudes, thus allowing that influence to become apparent in a manner that is less true of single, isolated interactions or judgments. (Shortly, we shall discuss the role of motivation and opportunity to control one's activated attitudes in more detail.) In any case, these findings illustrate that implicitly-measured attitudes can predict important, long-term outcomes. The finding regarding interracial roommate relationships, in particular, is difficult to reconcile with the idea that implicitly-

measured intergroup bias passes through the minds of individuals, but does not reside in them.

The Attitude-Nonattitude Continuum

This is not to say that the situation cannot be an important influence in shaping what is revealed by implicit measurements for some individuals. Indeed, it is useful to think about such situational forces and the wave metaphor (Who starts it? Who might prove resistant to it?) in the context of what has been referred to as the attitude-nonattitude continuum. In a model first proposed by Fazio, Chen, McDonel, and Sherman (1982), attitudes are viewed as associations between an object and an evaluation that are stored in memory. Like other associations, attitudes can vary in their associative strength, both within an individual for different attitude objects and across individuals for the same attitude object. This differential associative strength means that attitudes vary in their chronic accessibility from memory (Fazio, Sanbonmatsu, Powell, & Kardes, 1986). When an individual lacks an object-evaluation association in memory, they lie at the nonattitude end of the continuum with respect to that object. Movement toward the other end of the continuum is characterized by availability (i.e., the existence of an object-evaluation association) and increasing accessibility (i.e., associative strength) of object-evaluation associations in memory. At the higher end of the attitude continuum, the strength of the object-evaluation association

may be sufficient for the object to be attitude-evoking in the sense it can activate the attitude automatically when encountered by the individual (Fazio et al., 1986). Position along this attitude-nonattitude continuum has been found to determine the extent to which the attitude impacts attention, information processing, judgment, and behavior, thus influencing both attitude stability and attitude-behavior consistency (Fazio, 2007). Clearly, this conceptualization assigns significance to the very same notion of accessible links that Payne et al. emphasize.

Importantly, the model also offers an insight regarding the instability of implicit measures of intergroup bias. For individuals with relatively strong object-evaluation associations, implicit measures should capture chronically-accessible attitudes. It is for such individuals that attitudes should be automatically-activated in response to the primes that are presented in AMP or EPM. It is for such individuals that the specific nature of the response mapping imposed by the IAT should matter, matching or not matching their personal associations.

But what about an individual with a relatively weak object-evaluation association? The attitude-nonattitude continuum suggests that the individual's responses are more likely to be influenced by momentary situational influences. Indeed, a large body of literature indicates that when attitudes are relatively weak, responses to an attitude object become more sensitive to situational influence (see,

for example, Petty & Krosnick, 1995; Fazio, 2007). For example, for individuals that have relatively weak racial attitudes, either the immediate context or the larger-scale cultural context in which they are immersed may be responsible for shaping responses on implicit measures of intergroup bias to a significant extent. Such situational influence could account for the temporal instability of certain implicit measures, as the scores of those with relatively weak attitudes vary markedly because their responses are shaped by situational influences that themselves vary over time. But, any such influences should be tempered among those for whom the more chronically "accessible links" (to use Payne et al.'s term) are those reflecting their own attitudes.

At this point, it may be beneficial to return to the metaphor of the wave at a stadium to illustrate our theoretical position. When the average sports fan is in a stadium and someone starts the wave, they will most likely participate. That is, the situational context determines their behavior. We argue that this occurs because the average sports fan either has relevant supportive attitudes toward the team and the wave or does not have particularly strong attitudes in one direction or the other. However, we would predict that if a sports fan had a strong negative attitude toward the home team or toward the wave itself, they would not participate, independent of what was happening around them. This appears to be the case for San Francisco Giants fans, who have a distaste for the wave (Grossberg & Allen-Price, 2017). Usually, when someone tries to get the wave

started, other fans do not participate and the instigators are booed. So, if we wanted to predict whether a person or section will stand up when the wave is going around a stadium, knowing whether the wave is going around the stadium would be useful if the person or section had a weak relevant attitude, but if they held a strong negative attitude, the situation would matter to a much smaller extent. It would be their attitude (a stable, personal object-evaluation association) that would better predict their response, in this case their resistance to the wave.

Extrapersonal Associations

Importantly, various implicit measures may themselves be differentially sensitive to situational influence. Precisely this has been argued with respect to the IAT as it is traditionally implemented (Olson & Fazio, 2004). At its core, the traditional IAT assesses the ease with which individuals can pair a category with either positive or negative valence. Mentally pairing these concepts need not be purely a matter of one's attitudes. As noted in the target article, it has been suggested that the IAT, as traditionally administered, is open to contamination by *extrapersonal associaitions*–associations that are available in memory, but are not relevant to the expected outcome of personally interacting with the attitude object (i.e., nonattitudinal associations; Olson & Fazio, 2004). To remedy this problem, an IAT variant called the personalized IAT was developed. In this task, features of the traditional IAT that encouraged the use of extrapersonal associations were

altered. Specifically, the labels “Pleasant” and “Unpleasant” were changed to “I like” and “I don’t like,” error feedback was eliminated, and non-normatively positive or negative items were employed for the pleasant-unpleasant categorization task. In one experiment, individuals administered a traditional apple-candy IAT showed a preference for apples (consistent with normative views of apples as good and healthy and candy as bad and unhealthy), whereas those administered a personalized IAT showed no distinct preference. Research also has shown the personalized IAT to more effectively discriminate cigarette smokers from non-smokers, presumably because the traditional IAT is influenced by the normative negative view of smoking for both groups (Bardin, Perrissol, Launay, & Escoubes, 2014; De Houwer, Custers, & De Clercq, 2006). In another study concerning presidential candidates, the personalized IAT yielded stronger correspondence with measures of candidate preference, voting intentions, political party affiliation than did the traditional IAT (Olson & Fazio, 2004). Most importantly, when racial groups have been used as attitude objects, the personalized IAT has revealed lower levels of prejudice than the traditional IAT (Olson & Fazio, 2004), as well as stronger associations with a predictor of personal prejudice (Dambrun, Villate, & Richetin, 2008).

The susceptibility of the traditional IAT to situational influences was experimentally demonstrated in experiments by Han, Olson, and Fazio (2006). After having developed and rehearsed attitudes toward two novel stimuli, one of

which was clearly superior to the other, participants were exposed to a video interview in which two young boys offered comments that were consistent with the objectively correct value of those stimuli or comments that were clearly invalid. Despite this invalidity and despite judging the boys as irrational and foolish, participants in this latter condition exhibited a significantly reduced preference for the superior stimulus on a traditional IAT. These findings illustrate the very process that Payne et al. emphasized. However, additional findings from this research also merit highlighting. In particular, neither a personalized IAT nor an evaluative priming measure revealed any such attenuation in the preference for the superior or more preferable stimulus as a consequence of the situational influence manipulation. In another experiment, a contextual manipulation intended to momentarily alter the relative favorability of insects versus flowers (by exposure, in a presumably separate, earlier experiment, to a fictional narrative concerning a post-nuclear war environment in which insects were more favorable than flowers) attenuated the typical preference for flowers on a subsequent traditional IAT, but not on a personalized IAT (Han, Czellar, Olson, & Fazio, 2010). Thus, implicit measures themselves can vary in the extent to which they are affected by momentarily salient information.

As the work reviewed above suggests, this greater sensitivity to momentarily salient information may contribute to the traditional IAT's temporal instability and modest associations with behavior. For some individuals, the IAT

score is contaminated by extrapersonal associations that may themselves vary as a function of time and context. The personalized IAT appears to be relatively free of such contextual sensitivity. Initial work suggests that it is a useful tool to study intergroup bias (Dambrun et al., 2008; Olson & Fazio, 2004), but as mentioned in the target article, more research is needed in this area.

Motivation and Opportunity to Control Prejudiced Reactions

A component of the third puzzle discussed in the target article is the meta-analytic findings of a modest association, at the level of the individual, between scores on implicit measures of intergroup bias and discriminatory behavior. We argue that by focusing on simple correlations, these meta-analyses (and the individual investigations of which they consist) are insensitive to the complexity of the relation between attitudes and behavior. As mentioned earlier, attitude accessibility plays a moderating role in attitude-behavior consistency. Moreover, the complexity of that relation is also the result of an individual's motivation and opportunity (i.e., having sufficient time and resources) to control prejudiced reactions – factors that can intervene and drastically alter the nature of the attitude-behavior relation. In fact, the direction of this relation can even be reversed as a result of motivated overcorrection on the part of individuals who have negative attitudes automatically-activated by members of a given group, but

are both motivated to control their prejudiced reactions and have sufficient opportunity to do so.

Fazio et al. (1995) provided a demonstration of this complex relation between implicitly-measured racial attitudes and prejudicial responses. Participants in this study completed an individual difference measure of their motivation to control prejudiced reactions against Black individuals (Motivation to Control Prejudiced Reactions Scale; Dunton & Fazio, 1997), a measure of prejudicial responding against Black individuals called the Modern Racism Scale (MRS; McConahay, 1986), and a Black-White EPM. Their results revealed no significant main effect of implicitly-measured racial attitudes on scores on the MRS. Thus, at the level of a simple correlation, the association was indeed negligible. However, there was a significant interaction between implicitly-measured racial attitudes and motivation to control prejudiced reactions. For those relatively low in motivation, their implicitly-measured racial attitudes predicted responses on the MRS such that those with more negative attitudes displayed more prejudicial responses. Interestingly, among individuals relatively high in motivation, there was a trend such that more negative implicitly-measured attitudes predicted less prejudicial responses on the MRS, suggesting that such individuals engaged in motivated correction of their automatically-activated negativity. Similar interactive findings have been observed in numerous other studies (see Fazio & Olson, 2014, for a review).

Thus, the motivation to control prejudiced reactions can influence the extent to which implicitly-measured attitudes predict behavior. What may appear to be a modest or nonsignificant association between implicitly-measured attitudes and behavior may belie a more complicated relation.

Conclusion

The target article by Payne et al. tackled apparent puzzles in the implicitly-measured intergroup bias literature and outlined a thought-provoking model to account for them. Their model provides valuable insights, especially the possibility that variance due to systematic situational forces may be mistaken for error variance. And, we certainly applaud their suggestions that discrimination might be reduced effectively by re-structuring situations. Nonetheless, we argue that implicit measures of bias do not simply assess situational forces. Instead, we maintain that the individual occupies an important role in implicitly-measured intergroup bias. Understanding the relative influences of situations and individuals on implicit measures of bias, and the relation of the scores provided by the measures to behavior, requires that three things be kept in mind. First, the implicit measurement scores of individuals with relatively weak attitudes may be influenced by the situational context to a greater degree than is the case for individuals with stronger, more accessible attitudes. Thus, the strength of the individuals' attitudes should matter. Second, some implicit measures of

intergroup bias are more sensitive to the situational context than others. For example, the traditional IAT appears to be more open to the influence of momentarily salient extrapersonal associations than the personalized version. Thus, the specific implicit measure of bias that is employed matters. Third, attitudes do not exist in a vacuum. They interact with an individual's motivation to control their reactions and their opportunity to do so. Thus, motivation and opportunity to deliberate can alter the observed relation between implicit measures and subsequent judgments or behavior. We believe that these points shed light on the puzzles outlined in the target article without minimizing the role of the individual.

References

Bardin, B., Perrissol, J.P., Launay, C., & Escoubé. F. (2014). Personlized SC-IAT: A possible way of reducing the influence of societal views on assessments of implicit attitudes toward smoking. *Psychological Reports: Disability & Trauma, 115,* 13-25. doi: 10.2466/18.07PR0.115c10z8

Capers, Q., Clinchot, D., McDougle, L., & Greenwald, A. G. (2017). Implicit racial bias in medical school admissions. *Academic Medicine: Journal of the Association of American Medical Colleges, 92,* 365-369. doi: 10.1097/ACM.0000000000001388

Castelli, L., Carraro, L., Gawronski, B., & Gava, K. (2010). On the determinants of implicit evaluations: When the present weighs more than the past. *Journal of Experimental Social Psychology, 46,* 186-191. doi: 10.1016/j.jesp.2009.10.006

Dambrun, M., Villate. M., & Richetin, J. (2008). Implicit racial attitudes and their relationships with explicit personal and cultural beliefs: What personalized and traditional IATs measure. *Current Research in Social Psychology, 13,* 185-198.

De Houwer, J., Custers, R., & De Clerq, A. (2006). Do smokers have a negative implicit attitude toward smoking? *Cognition & Emotion, 20,* 8, 1274-1284. doi: 10.1080/02699930500484506

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. doi: 10.1177/0146167297233009

Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition*, 25, 603-637.

Fazio, R. H., Chen, J., McDonel, E. C., & Sherman, S. J. (1982). Attitude accessibility, attitude-behavior consistency and the strength of the object-evaluation association. *Journal of Experimental Social Psychology,* 18, 339-357. doi: 10.1016/0022-1031(82)90058-0

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. doi: 10.1037/0022-3514.69.6.1013

Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition research: Their meaning and use. *Annual Review of Psychology, 54*, 297-327. doi: 10.1146/annurev.psych.54.101601.145225

Fazio, R. H., & Olson, M. A. (2014). The MODE model: Attitude-behavior processes as a function of motivation and opportunity. In J. W. Sherman, B. Gawronski, & Y. Trope (Eds.), *Dual process theories of the social mind* (pp. 155-171). New York: Guilford Press.

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. doi: 10.1037/0022-3514.50.2.229

Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74,* 1464–1480. doi:10.1037/0022-3514.74.6.1464

Greenwald, A. G., & Nosek, B.A. (2008). Attitudinal dissociation: What does it mean? In Petty, R. E., Fazio, R. H., & Briñol, P. (Eds.), *Attitudes: Insights from the new implicit measures* (pp. 65–82). Hillsdale, NJ: Lawrence Erlbaum Associates.

Grossberg, A., & Allen-Price, A. G. (2017, April 13). *Why do Giants fans hate the wave? Oakland started it.* Retrieved from https://ww2.kqed.org/news/2015/07/21/the-wave-was-born-in-oakland-and-some-giants-fans-want-it-dead/

Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of implicit attitudes. *Journal of Experimental Psychology General*, *143*, 1369–1392. doi: 10.1037/a0035028

Han, H.A., Czellar, S., Olson, M.A., & Fazio, R.H. (2010). Malleability of attitudes or malleability of the IAT? *Journal of Experimental Social Psychology, 46,* 286-298. doi: 10.1016/j.jesp.2009.11.011

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. doi: 10.1016/j.jesp.2005.04.006

Kang, J., Bennett, M., Carbado, D., Casey, P., Dasgupta, N., Faigman, D. L., Godsil, R., ... Mnookin, J. (2012). Implicit bias in the courtroom. *UCLA Law Review, 59,* 5, 1124-1186.

Lee, S., Rogge, R. D., & Reis, H. T. (2010). Assessing the seeds of relationship decay: Using implicit evaluations to detect the early stages of disillusionment. *Psychological Science, 21*(6), 857 -864. doi: 10.1177/0956797610371342

McConahay, J.B. (1986). Modern racism, ambivalence, and the Modern Racism Scale. In Dovidio, J.F., & Gaertner, S.L. (Eds.), *Prejudice, discrimination, and racism* (pp. 91-125). San Diego, CA: Academic Press.

McNulty, J. K., Olson, M. A., Meltzer, A. L., & Shaffer, M. J. (2013). Though they may be unaware, newlyweds implicitly know whether their marriage will be satisfying. *Science, 342*, 1119-1120. doi: 10.1126/science.1243140

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. doi: 10.1037/0022-3514.86.5.653

Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. D. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology, 89,* 277–293. doi:10.1037/0022-3514.89.3.277

Petty, R.E. & Krosnick, J.A. (Eds.). (1995). *Attitude Strength: Antecedents and Consequences*. Mahwah, NJ: Lawrence Erlbaum.

Ranganath, K. A., Smith, C. T., & Nosek, B. A. (2008). Distinguishing automatic and controlled components of attitudes from direct and indirect measurement methods. *Journal of Experimental Social Psychology, 44,* 386–396. doi: 10.1016/j.jesp.2006.12.008

Rudman, L. A. (2004). Sources of Implicit Attitudes. *Current Directions in Psychological Science, 13,* 2, 79-82. doi: 10.1111/j.0963-7214.2004.00279.x

Towles-Schwen, T., & Fazio, R. H. (2001). On the origins of racial attitudes: Correlates of childhood experiences. *Personality and Social Psychology Bulletin,* 27, 162-175. doi: 10.1177/0146167201272003

Towles-Schwen, T., & Fazio, R. H. (2006). Automatically-activated racial attitudes as predictors of the success of interracial roommate relationships. *Journal of Experimental Social Psychology*, 42, 698-705. doi: 10.1016/j.jesp.2005.11.003