

Implicit Attitude Measures

GREGORY MITCHELL and PHILIP E. TETLOCK

Abstract

Owing to concerns about the willingness and ability of people to report their attitudes accurately in response to direct inquiries, psychologists have developed a number of unobtrusive, or implicit, measures of attitudes. The most popular contemporary implicit measures equate spontaneous responses to stimuli with attitudes about those stimuli. Although these measures have been used to open important new lines of inquiry, they suffer from reliability and construct validity problems and administration limitations. Researchers conducting basic research on attitudes may fruitfully utilize implicit measures as part of a multipronged measurement strategy, but researchers seeking to predict behavior from attitudes should continue to rely on explicit measures of attitudes, taking care to minimize reactive bias and to formulate the attitude questions at the same level of specificity as the behavior to be predicted.

INTRODUCTION

Is tennis more enjoyable than golf? Should same-sex couples be permitted to adopt children? We have little reason to suspect that social norms will lead to deceptive responses to the first question, but many people may be unwilling to answer the second question honestly for fear of offending others or being perceived as intolerant. If some strategic gain is to be had from favoring golf over tennis, such as ingratiation of a superior at work, then impression management goals may cause insincere responses even to the first question (Tedeschi, Schlenker, & Bonoma, 1971). Allowing anonymous responses to both questions, if permitted by the research design, may alleviate concerns that the context will influence the responses, but we must still worry whether individuals *can* give honest answers to these questions given research demonstrating disparities between stated and behaviorally-express preferences (e.g., Nisbett & Wilson, 1977). These concerns—about reactive bias arising from social desirability pressures or from the related but situation-specific problem of impression management and about the lack of reliable access to one's own preferences

through conscious deliberation—gave rise to efforts to develop unobtrusive measures of attitudes.

Today a variety of measures exist for measuring unobtrusively, or “implicitly,” an individual’s evaluative stance toward political, social, economic, and personal matters. The most popular new measures examine how fast an attitude object can be categorized positively or negatively and attach significance to millisecond differences in response times (e.g., Greenwald, McGhee, & Schwartz, 1998), with less than a second often separating positive from negative attitude ascriptions. These measures define attitudes as evaluative associations with an attitude object and require no conscious endorsement of the evaluation or behavioral manifestation for an attitude to be ascribed to an individual. Any gains in nonreactivity and access to unmediated thought obtained through this measurement approach come with serious questions about the reliability, construct validity, and predictive validity of these new measures. Until these issues are sorted out, these new measures are most appropriate for basic attitude research rather than as an alternative to traditional explicit measures of attitudes in research where a measure of attitudes is needed as one component of the project. For instance, implicit attitude measures may be useful in exploring the underlying psychological components of consumer preferences, but surveys that explicitly question consumers about their product preferences and purchase intentions are likely to be more predictive of purchasing behavior (Greenwald, Poehlman, Uhlmann, & Banaji, 2009) and much easier to use.

FOUNDATIONAL AND CUTTING-EDGE RESEARCH

Concerns about inaccurate responses to interview and survey questions have perpetually dogged social scientists (Crosby, Bromley, & Saxe, 1980; Ostrom, 1973). In 1966, the methodology experts Webb, Campbell, Schwartz, and Sechrest (1966) devoted an entire book to unobtrusive measures of attitudes and other psychological phenomena, providing a survey and analysis of a wide range of observational, archival, and physical-trace methods that continue to be used to measure unobtrusively what people think and feel about various topics. Beginning in the earliest days of attitude research, psychologists embarked on a quest to find a measure of attitudes that does not rely on participant introspection and honesty (Vargas, Sekaquaptewa, & von Hippel, 2007). The long journey continues.

Initial attempts by social psychologists to overcome the limits of self-report-based, or “explicit,” measures of attitudes relied on stealth. In one particularly influential approach, Jones and Sigall (1971) employed what they called the “bogus pipeline” to attitudes: after connecting participants to a device that supposedly measures attitudinal direction and intensity

using sensitive physiological measurements, participants must estimate their feelings toward various attitude objects for comparison with their “true” feelings as measured by the device. The key assumption behind the bogus pipeline paradigm is that participants will be motivated to give truthful self-reports when faced with the prospect of contradiction by the sophisticated measuring device that supposedly provides a pipeline to the attitudinal soul. Although the bogus pipeline procedure produced reliable effects that seemed to be less contaminated by social-desirability bias (e.g., in studies of racial attitudes, on average participants hooked up to the bogus pipeline machine reported greater prejudice than participants in the control condition who completed traditional explicit measures of attitudes), ethical concerns, construct validity questions, and technological changes led to greatly reduced use of the bogus pipeline paradigm within just two decades of its introduction (Roese & Jamieson, 1993).

In the 1980s, psychologists began measuring the direction and strength of attitudes by measuring the speed with which attitude objects are paired with negative or positive evaluative terms. These new methods took advantage of technological innovations that allowed researchers to present many kinds of stimuli for very brief periods of time, via computer terminals, and measure response times with great sensitivity, also via computer terminals. By presenting stimuli at subliminal or just supraliminal levels and requiring quick responses, these tasks are thought to limit the influence of strategic responding (it is standard with these measures to exclude responses that exceed some temporal threshold above which responding is deemed deliberate rather than spontaneous). The key assumptions behind this approach are that (i) stronger associations between evaluative and attitude-object categories will produce shorter response times on speeded tasks in which stimuli from the evaluative and attitude-object categories must be compared, (ii) “attitudes” do not require access to intentional-level responding or declarative memory (i.e., deliberate endorsement of an evaluation of an attitude object is not a necessary element of an attitude), and (iii) quick, spontaneous responses reveal automatic, or relatively unconscious, associations among the evaluative and attitude-object categories.

Fazio, Sanbonmatsu, Powell, and Kardes (1986) introduced the first of these new measures that rely on spontaneous responses to attitude objects to assess attitudes, an approach Fazio, Jackson, Dunton, and Williams (1995) later suggested could be a bona fide pipeline to our true attitudes. Fazio and colleagues’ procedure, which has come to be known as evaluative, affective, or sequential priming, involves multiple trials in which participants briefly see the name of an attitude object (e.g., snake) followed by a positive or negative adjective (e.g., scary); on each trial, participants must categorize the adjective term as positive or negative as quickly as possible. If responses to the

negative adjectives are faster than responses to the positive adjectives, then the attitude object is said to facilitate negative responding, and this facilitation is taken as evidence of negative associations with the attitude object and consequently evidence of a negative attitude toward the object.¹ This automaticity-based approach spawned a number of similar measures, and these automaticity-based measures now dominate attitude research within social psychology.²

In 1998, Greenwald *et al.* (1998) introduced what has become the most popular implicit measure of attitudes, the Implicit Association Test (IAT). The IAT presents participants with brief images of stimuli to be classified as quickly as possible over many trials; the stimuli consist of two sets of stimuli that serve as attitude objects and positive and negative adjectives. The two groups of attitude-object stimuli are paired, over successive trials, with either the positive or negative adjectives, and in each trial the participant is asked to press one computer key if one type of attitude object or one type of adjective is observed and to press a different computer key if the other type of attitude object or the other type of adjective is observed. For instance, on the racial attitudes IAT, in one block of trials participants must tap a left-hand key on the computer if an image of a white face or a positive word is shown and a right-hand key if an image of a black face or negative word is shown, and on another block of trials white faces share a response key with negative words while black faces share a response key with positive words. If response times are faster in the first set of trials relative to the second set of trials, then the participant is said to hold a more positive attitude toward whites relative to blacks. The assumption is that a congruence of associations between the attitude object and words of a particular valence facilitates classification on the trials where the attitude object and words of that valence share a response key (e.g., persons holding positive associations with the white race should find it easier to classify white faces/positive terms than white faces/negative terms).

The IAT requires that stimuli from two attitude-object categories be placed into opposition because the focal measurement outcome is a difference score: the average response time when one set of attitude-object stimuli

1. In 1997, Wittenbrink, Judd, and Park (1997) introduced a very similar procedure that has come to be known as semantic priming, but this procedure is used primarily to assess stereotypes as opposed to attitudes. In this procedure, participants see words from a target category (e.g., black or white in a study of racial stereotypes) followed by meaningful or meaningless letter strings (e.g., possible trait terms or nonsense words), and participants must decide as quickly as they can whether the letters formed a word or not. If the target category term facilitates responses to positive or negative trait terms, then the participants is assumed associate positive or negative stereotypes with the target category.

2. Some new implicit attitude measures do not measure reaction times but do try to take advantage of spontaneous responses to stimuli (see Vargas *et al.*, 2007). For instance, Isen, Labroo, and Durlach (2004) exposed participants to attitude objects and then asked the participants to fill in the blanks on words that could be completed to have positive or negative meaning. The valence of the completed words was taken as an indication of whether the attitude object primed positive or negative associations.

is paired with positive terms and the other set of attitude-object stimuli is paired with negative terms minus the average response time when the pairings are reversed.³ The inherently relativistic nature of the IAT leads to interpretation problems (e.g., a difference in response times on the racial attitudes IAT may reflect greater negativity toward blacks or greater positivity toward whites, and persons with similar scores on the IAT may hold very different patterns of associations with the attitude objects) and prompted the creation of similar measures that examine only one attitude object at a time. In the Go/No-Go Association Task (GNAT) (Nosek & Banaji, 2001), participants see stimuli from the attitude-object category and positive or negative terms and distractor terms over multiple sets of trials; on one set of trials, participants press a computer key if a member of the attitude-object category or a positive term is viewed (the go response) and do nothing if a distractor stimuli is viewed (the no-go response), and on another set of trials the go response applies to the attitude stimuli and negative terms. If greater sensitivity is shown when the attitude object is paired with positive terms, then the participant is said to hold a positive attitude toward the object; if greater sensitivity is shown when the attitude object is paired with negative terms, then the participant is said to hold a negative attitude. In the Extrinsic Affective Simon Task (EAST) (De Houwer, 2003), participants view stimuli from an attitude-object category in fonts of one of two colors and view positive or negative words in a white font. If a stimulus is presented in white font, the participant must classify the stimulus by its valence, and if the stimulus is in another color, then it must be classified by color. The assumption is that faster or more accurate responses when the attitude object is paired with the positive or negative terms indicates, respectively, either positive or negative attitude toward the attitude object. The EAST continues to be used, but De Houwer concluded that the IAT outperforms the EAST as a measure of attitudes (De Houwer & De Bruycker, 2007).

In 2005, Payne, Cheng, Govorun, and Stewart (2005) introduced the affect misattribution procedure (AMP), which, like the IAT, quickly became popular among psychologists, but unlike the IAT can be used to measure attitudes toward a single attitude object (i.e., it is not inherently relativistic in nature). In the AMP, participants are very briefly shown stimuli from the attitude-object category followed by a character from the Chinese alphabet and are asked whether the Chinese character is or more or less visually pleasant than the average Chinese character. If evaluations of Chinese characters tend to be positive after the attitude primes, then the participant is said to hold a positive attitude toward the attitude object, and negative attitude

3. For a full discussion of the algorithm presently used to score the IAT, see Greenwald, Nosek, and Banaji (2003).

ascriptions follow from negative evaluations of the Chinese characters on the heels of the attitude primes. The assumption is that people project their evaluations of the attitude prime onto the ambiguous Chinese symbols.

The latest frontier in implicit attitude measurement involves sophisticated physiological measures. Although physiological measures have long been used as indirect measures of attitudes (e.g., activity in facial muscles associated with positive reactions to stimuli has been taken to signify positive attitudes; Cacioppo & Petty, 1979), the latest physiological approaches employ functional brain imagery to monitor activation in areas of the brain thought to signify affective processing of the attitude objects presented to participants (see Ito & Cacioppo, 2007). Presently these approaches cannot be used outside a laboratory setting and can be applied only to small numbers of respondents, rendering them useful for basic research that seeks to examine the neurological basis of, or mechanisms underlying, attitudes but not for other types of research.⁴

OPEN QUESTIONS AND INSTRUMENT LIMITATIONS

The new implicit attitude measures, particularly the IAT, enjoy incredible popularity. Hundreds of IAT studies have been published since the IAT's introduction in 1998, the IAT has been adapted to measure a wide range of attitudes and stereotypes, and the popular press has embraced findings from the IAT research program (e.g., Gladwell, 2005). Popularity should not be mistaken for utility and validity. Although the IAT has advantages over some of its competitors, such as greater reliability, the popularity of the IAT appears to derive primarily from its adaptability, public dissemination of the programming code that makes the creation of new IATs relatively easy, and the tantalizing possibility that the IAT provides a pipeline to the unconscious that reveals deep-seated attitudes that many individuals did not even know they possessed. When considering whether to incorporate the IAT or another contemporary implicit attitude measure into a research project, social scientists should consider their limitations and the many open questions that surround these new implicit measures.⁵

4. Another set of implicit measures infer attitudes from a participant's approach or avoidance behavior in response to an attitude object, as measured, for example by pulling or pushing a lever (e.g., Chen & Bargh, 1999). We do not focus on these measures because they are much less popular presently than reaction-time-based measures and because of recent questions about what drives the approach and avoidance behavior observed in these tasks (see Gawronski, 2009).

5. We have observed the unfortunate tendency of researchers to treat the IAT as if it were a Likert scale that can be easily adapted to any study to measure attitudes without first engaging in the validation work needed to ensure that the attitude-object stimuli do not bias the results (see Nosek, Greenwald & Banaji, 2007) and without considering the limitations of the IAT, particular those arising from its relativistic approach to attitude measurement (see, e.g., Blanton *et al.*, 2007; Blanton & Jaccard, 2006).

First and foremost, the new implicit attitude measures can be difficult to implement. In a laboratory setting, the new implicit measures involve considerable time and effort, often requiring their own experimental session because of the instrumentation and multiple trials and involved, and it is not feasible to use some of the measures outside the laboratory. The automaticity-based implicit measures can be incorporated into online survey research, but the added time and effort required to complete these measures may tax respondents and lead to attrition and use of these measures comes at the cost of omitting alternative questions and measures (for a discussion of problems that may be encountered when seeking to incorporate implicit attitude measures into computer-based survey research, see Krosnick & Lupia, 2008).

Second, a fundamental concern of automaticity-based implicit measures is that participants not be able to mediate consciously their responses. Unfortunately, there is evidence that responses on implicit tasks are not beyond the control of respondents. Participants often infer the purpose behind the task and can intentionally alter their pattern of responses and thus the attitudes ascribed to them, controlled processes contribute to responses on the implicit tasks even when those processes are beyond the awareness of participants, and reactivity biases can affect responses on these measures (see, e.g., Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Czellar, 2006; Fiedler, Messner, & Bluemke, 2006; Frantz, Cuddy, Burnett, Ray, & Hart, 2004; Gawronski, 2009).

Even if a respondent is not aware of the purpose behind the implicit task or cannot consciously mediate her response, the person's observed behavior may be caused by something other than attitudes. This possibility gives rise to the third important open question for the new implicit measures: to what extent are responses on these measures contaminated by artifacts, such as individual differences in working memory that affect the speed with which information is processed? These new measures are not "process pure": they do not measure only the target construct of interest and some artifacts may significantly affect the measures taken by the new implicit measures (Nosek & Smyth, 2007). Factors such as the respondent's amount of practice on the task, age, general processing speed and ability to switch tasks quickly and effectively, and familiarity with the attitude-object stimuli, if not accounted for, will contaminate the results and lead to erroneous conclusions about the attitudes of respondents (see Blanton, Jaccard, Christie, & Gonzales, 2007; Mitchell & Tetlock, 2006).

More generally, the new implicit measures raise basic construct validity questions concerning the meaning of an attitude and how to go about measuring attitudes. The proper definition and operationalization of the attitude construct is beyond the scope of this essay, particularly given the long history

of debate over the attitude concept and the multiplicity of definitions offered (McGuire, 1985). But a researcher considering the use of the new implicit measures should be aware of ongoing debates about the proper definition of attitude and whether the new implicit measures actually measure anything that should be called an attitude. One prominent debate, engaging the inventors of evaluative priming and the IAT, concerns whether the IAT measures personal attitudes or cultural knowledge that should not be deemed a personal attitude (see Olson & Fazio, 2009; see also Arkes & Tetlock, 2004). One resolution of these definitional debates involves splitting the attitude construct in two: implicit measures tap into *implicit* attitudes, whereas explicit measures tap into *explicit* attitudes (e.g., Wilson, Lindsey, & Schooler, 2000). This compromise seeks to make sense of data showing that the measurements made by implicit and explicit measures sometimes converge and sometimes diverge by specifying the conditions under which, and the types of attitude objects for which, expressions of implicit attitudes are likely to depart from expressions of explicit attitudes (e.g., Hofmann, Gawronski, Gschwendner, Huy, & Schmitt, 2005; Nosek, 2005, 2007; Smith & Nosek, 2011). This body of research should be consulted before incorporating an implicit measure into a research project, lest one use an implicit measure for a situation or attitude object where no divergence is expected and thus incur unnecessary costs of using the implicit measure. We view labeling any association an attitude as too sweepingly reductionist an approach, which leads, among other things, to conflating things we believe with things that we suspect others believe and conflating objective observations with personal attitudes (such as recognizing the success of the Boston Red Sox versus having a positive attitude toward the Red Sox) (see also Petty, Briñol, & DeMarree, 2007). Nonetheless, some psychologists seem to embrace that very idea (see, e.g., Banaji, Nosek, & Greenwald, 2004).

Yet another problem with automaticity-based implicit measures is that they often exhibit low split-half and test-retest reliability scores (Fazio & Olson, 2003; Nosek, Greenwald, & Banaji, 2007). The IAT tends to outperform the evaluative priming procedure, though the IAT's test-retest reliability and internal consistency as measured by split-half reliability are both less than desired for measures of attitudes, which are supposed to be reasonably stable dispositions toward objects.⁶ Early tests with the AMP suggest that its reliability is comparable to that of the IAT (e.g., Payne, Govorun, & Arbucke, 2008).

6. These reliability estimates do not reflect the impact of systematic variations in the testing environment, which have also been shown to affect scores on implicit tasks, suggesting that the implicit measures assess transient states rather than stable associative networks (Mitchell & Tetlock, 2006; Smith & Conrey, 2007).

Finally, the new implicit measures often fail to outperform simple explicit measures of attitudes in the prediction of behavior.⁷ This finding should not be surprising given the fairly low reliability of the new measures, because low predictive validity follows from low reliability, nor given uncertainty about what exactly the new implicit measures measure. Greenwald *et al.* (2009) reported that across a number of domains explicit attitude measures performed better than, or as well as the IAT, including on sensitive topics concerning drug use, self-injury, and gender attitudes, but they reported that the IAT outperformed explicit measures when predicting behavior toward racial and other minority groups. However, Oswald, Mitchell, Blanton, Jaccard, & Tetlock (2013) performed a follow-up meta-analysis of the studies in which racial and ethnic attitude IATs were used to predict behavior and found that the IAT was a poor predictor of all types of behavior and was outperformed by even very simple explicit attitude measures. Cameron, Brown-Iannuzzi, and Payne (2012) conducted a meta-analysis of studies in which sequential priming measures were used to predict behavior and found that the priming measure and explicit measures did not significantly differ in their predictive validity. It appears that if steps are taken to minimize reactivity bias in response to explicit attitude measures (see Bradburn, Sudman, & Wansink, 2004; Tourangeau & Yan, 2007), and if the attitude queries are framed at the same level of specificity as the behavior to be predicted [as contemporary research into attitude-behavior relations counsels in order to increase predictive validity (see Oswald *et al.*, 2013)], then explicit attitude measures will provide equal or better prediction and be much simpler to implement than automaticity-based implicit measures.

CONCLUSION

If one is conducting basic or exploratory research on attitudes, then incorporating an implicit attitude measure into the research may be worthwhile. However, the latest incarnations of implicit measures of attitudes, which emphasize automatic responses to stimuli, are not good candidates for addition to studies where the goal is to obtain a reliable and predictive measure of attitudes or where attitudes are being assessed outside the laboratory. The latest implicit attitude measures do not provide efficient approaches to

7. A related problem for the new implicit measures concerns a lack of discrimination among respondents. The racial attitudes IAT, for instance, leads to many inaccurate predictions about how respondents will behave in the presence of minorities (Fiedler *et al.*, 2006; Mitchell & Tetlock, 2006). With socially sensitive matters, such as the ascription of prejudicial attitudes to persons, and with economic matters, such as the prediction of product preferences in consumer product research, this inability to discriminate can have serious consequences for both respondents and researchers. Furthermore, because outliers may drive observed correlations between implicit attitudes and behavior (Blanton *et al.*, 2009), researchers should not assume constant relationships between scores on implicit measures and behavioral variables.

gathering attitudinal data for a host of reasons: The new measures undergo serious reliability and construct validity problems, can suffer from reactive bias just as explicit attitude measures can, rarely outperform explicit measures of attitudes with respect to behavioral prediction, and many of the new measures are difficult and time-consuming to implement. Explicit measures of attitudes are much easier to use, reactive bias associated with explicit measures can be minimized and monitored, and explicit measures will likely provide equal or better predictive validity than the latest generation of implicit attitude measures. The current popularity of implicit attitude measures appears to be driven more by their availability and novelty, and the never-ending quest by social psychologists to find a bona fide pipeline to “true” attitudes, than by the scientifically demonstrated validity and utility of the new measures.

REFERENCES

- Arkes, H., & 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., Nosek, B. A., & Greenwald, A. G. (2004). No place for nostalgia in science: A response to Arkes & Tetlock. *Psychological Inquiry, 15*, 279–289.
- Blanton, H., & Jaccard, J. (2006). Arbitrary metrics in psychology. *American Psychologist, 61*, 27–41.
- Blanton, H., Jaccard, J., Christie, C., & Gonzales, P. M. (2007). Plausible assumptions, questionable assumptions and post hoc rationalizations: Will the real IAT please stand up? *Journal of Experimental Social Psychology, 43*, 393–403.
- Blanton, H., Jaccard, J., Klick, J., Mellers, B., Mitchell, G., & Tetlock, P. E. (2009). Strong claims and weak evidence: Reassessing the predictive validity of the IAT. *Journal of Applied Psychology, 94*, 567–582. doi:10.1037/a0014665
- Bradburn, N., Sudman, S., & Wansink, B. (2004). *Asking questions: The definitive guide to questionnaire design—for market research, political polls, and social and health questionnaires*. San Francisco, CA: Josey-Bass.
- Cacioppo, J. T., & Petty, R. E. (1979). Attitudes and cognitive response: An electrophysiological approach. *Journal of Personality and Social Psychology, 37*, 2181–2199.
- Cameron, C. D., Brown-Iannuzzi, J., & Payne, B. K. (2012). Sequential priming measures of implicit social cognition: A meta-analysis of associations with behaviors and explicit attitudes. *Personality and Social Psychology Review, 16*, 330–350.
- Chen, M., & Bargh, J. A. (1999). Nonconscious approach and avoidance behavioral consequences of the automatic evaluation effect. *Personality and Social Psychology Bulletin, 25*, 215–224.
- Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology, 89*, 469–487. doi:10.1037/0022-3514.89.4.469

- Crosby, F., Bromley, S., & Saxe, L. (1980). Recent unobtrusive studies of black and white discrimination and prejudice: A literature review. *Psychological Bulletin*, *87*, 546–563.
- Czellar, S. (2006). Self-presentational effects in the Implicit Association Test. *Journal of Consumer Research*, *16*, 92–100.
- De Houwer, J. (2003). The extrinsic affective Simon task. *Experimental Psychology*, *50*, 77–85.
- De Houwer, J., & De Bruycker, E. (2007). The implicit association test outperforms the extrinsic affective Simon task as an implicit measure of inter-individual differences in attitudes. *British Journal of Social Psychology*, *46*, 401–421.
- Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027.
- Fazio, R. H., & Olson, M. A. (2003). Implicit measures in social cognition: Their meaning and use. *Annual Review of Psychology*, *54*, 297–327.
- Fazio, R. H., Sanbonmatsu, D. M., Powell, M. C., & Kardes, F. R. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, *50*, 229–238.
- Fiedler, K., Messner, C., & Bluemke, M. (2006). Unresolved problems with the “I”, the “A”, and the “T”: A logical and psychometric critique of the Implicit Association Test (IAT). *European Review of Social Psychology*, *17*, 74–147.
- Frantz, C., Cuddy, A. J. C., Burnett, M., Ray, H., & Hart, A. (2004). A threat in the computer: The race Implicit Association Test as a stereotype threat experience. *Personality and Social Psychology Bulletin*, *30*, 1611–1624.
- Gawronski, B. (2009). Ten frequently asked questions about implicit measures and their frequently supposed, but not entirely correct answers. *Canadian Psychology*, *50*, 141–150.
- Gladwell, M. (2005). *Blink*. New York, NY: Little, Brown and Company.
- 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., Nosek, B. A., & Banaji, M. R. (2003). Understanding and Using the Implicit Association Test: I. An Improved Scoring Algorithm. *Journal of Personality and Social Psychology*, *85*, 197–216.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97*, 17–41.
- Hofmann, W., Gawronski, B., Gschwendner, T., Huy, L., & Schmitt, M. (2005). A meta-analysis on the correlation between the implicit association test and explicit self-report measures. *Personality and Social Psychology Bulletin*, *31*, 1369–1385.
- Isen, A. M., Labroo, A. A., & Durlach, P. (2004). An influence of product and brand name on positive affect: Implicit and explicit measures. *Motivation and Emotion*, *28*, 43–63.
- Ito, T. A., & Cacioppo, J. T. (2007). Attitudes as mental and neural states of readiness: Using physiological measures to study implicit attitudes. In B. Wittenbrink & N.

- Schwarz (Eds.), *Implicit measures of attitudes* (pp. 125–158). New York, NY: Guilford Press.
- Jones, E. E., & Sigall, H. (1971). The bogus pipeline: A new paradigm for measuring affect and attitude. *Psychological Bulletin*, *76*, 349–364.
- Krosnick, J. A. & Lupia, A. (2008). Decisions made about implicit attitude measurement in the 2008 American National Election Studies. Memorandum. Retrieved from http://www.electionstudies.org/announce/newsltr/20090625_IAT.pdf
- McGuire, W. J. (1985). Attitudes and attitude change. In G. Lindzey & E. Aronson (Eds.), *Handbook of social psychology* (Vol. 2, pp. 233–346). New York, NY: Random House.
- Mitchell, G., & Tetlock, P. E. (2006). Antidiscrimination law and the perils of mindreading. *Ohio State Law Journal*, *67*, 1023–1121.
- 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. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, *134*, 565–584.
- Nosek, B. A. (2007). Implicit-explicit relations. *Current Directions in Psychological Science*, *16*, 65–69.
- Nosek, B. A., & Banaji, M. R. (2001). The go/no-go association task. *Social Cognition*, *19*, 161–176.
- Nosek, B. A., Greenwald, A. G., & Banaji, M. R. (2007). The Implicit Association Test at age 7: A methodological and conceptual review. In J. A. Bargh (Ed.), *Automatic processes in social thinking and behavior* (pp. 265–292). New York, NY: Psychology Press.
- Nosek, B. A., & Smyth, F. L. (2007). A multitrait-multimethod validation of the Implicit Association Test: Implicit and explicit attitudes are related but distinct constructs. *Experimental Psychology*, *54*, 14–29.
- Olson, M. A., & Fazio, R. H. (2009). Implicit and explicit measures of attitudes: The perspective of the MODE model. In R. E. Petty, R. H. Fazio & P. Briñol (Eds.), *Attitudes: Insights from the new implicit measures* (pp. 19–64). New York, NY: Psychology Press.
- Ostrom, T. M. (1973). The bogus pipeline: A new *ignis fatuus*? *Psychological Bulletin*, *79*, 252–259.
- Oswald, F., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology*, *105*(2), 171–192.
- Payne, B. K., Cheng, C. M., Govorun, O., & Stewart, B. (2005). An inkblot for attitudes: Affect misattribution as implicit measurement. *Journal of Personality and Social Psychology*, *89*, 277–293.
- Payne, B. K., Govorun, O., & Arbuckle, N. L. (2008). Automatic attitudes and alcohol: Does implicit liking predict drinking? *Cognition and Emotion*, *22*, 238–271.
- 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.
- Roese, N. J., & Jamieson, D. W. (1993). Twenty years of bogus pipeline research: A critical review and meta-analysis. *Psychological Bulletin*, *114*, 363–375.

- Smith, E. R., & Conrey, F. R. (2007). Mental representations are states, not things: Implications for implicit and explicit measurement. In B. Wittenbrink & N. Schwarz (Eds.), *Implicit measures of attitudes* (pp. 247–264). New York, NY: Guilford Press.
- Smith, C. T., & Nosek, B. A. (2011). Affective focus increases the concordance between implicit and explicit attitudes. *Social Psychology*, *42*, 300–313.
- Tedeschi, J. T., Schlenker, B. R., & Bonoma, T. V. (1971). Cognitive dissonance: Private ratiocination or public spectacle? *American Psychologist*, *26*, 685–695.
- Tourangeau, R., & Yan, T. (2007). Sensitive questions in surveys. *Psychological Bulletin*, *133*, 859–883.
- Vargas, P. T., Sekaquaptewa, D., & von Hippel, W. (2007). Armed only with paper and pencil: “Low-tech” measures of implicit attitudes. In B. Wittenbrink & N. Schwarz (Eds.), *Implicit measures of attitudes* (pp. 125–158). New York, NY: The Guilford Press.
- Webb, E. J., Campbell, D. T., Schwartz, R. D., & Sechrest, L. (1966). *Unobtrusive measures: Nonreactive research in the social sciences*. Chicago, IL: Rand McNally.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, *107*, 101–126.
- Wittenbrink, B., Judd, C. M., & Park, B. (1997). Evidence for racial prejudice at the implicit level and its relationship with questionnaire measures. *Journal of Personality and Social Psychology*, *72*, 262–274.

FURTHER READING

- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, *135*, 347–368.
- Fishbein, M., & Ajzen, I. (2010). *Predicting and changing behavior: The reasoned action approach*. New York, NY: Psychology Press.
- Gawronski, B., & Payne, B. K. (Eds.) (2010). *Handbook of implicit social cognition: Measurement, theory, and applications*. New York, NY: Guilford Press.
- Petty, R. E., Fazio, R. H., & Briñol, P. (Eds.) (2009). *Attitudes: Insights from the new implicit measures*. New York, NY: Psychology Press.
- Wilson, T. D., & Dunn, E. (2004). Self-knowledge: Its limits, value, and potential for improvement. *Annual Review of Psychology*, *55*, 493–518.
- Wittenbrink, B., & Schwarz, N. (Eds.) (2007). *Implicit measures of attitudes*. New York, NY: Guilford Press.

GREGORY MITCHELL SHORT BIOGRAPHY

Gregory Mitchell is the Joseph Weintraub-Bank of America Distinguished Professor of Law and Thomas F. Bergin Teaching Professor of Law at the University of Virginia. Mitchell, who holds a JD and a PhD in psychology, writes on intergroup relations, rational choice, social scientific methodology, and the application of social science to public policy issues.

Personal webpage: <http://www.law.virginia.edu/lawweb/Faculty.nsf/FHPbI/1191856>

Curriculum vitae: http://www.law.virginia.edu/pdf/faculty/mitchell_cv.pdf

PHILIP E. TETLOCK SHORT BIOGRAPHY

Philip E. Tetlock is the Leonore Annenberg University Professor in Democracy and Citizenship at the University of Pennsylvania. Tetlock studies judgment and decision-making, expert prediction, and intergroup relations. Tetlock has edited a number of books on social science topics and wrote *Expert Political Judgment: How Good Is It? How Can We Know?* (2006), which was awarded the University of Louisville Grawemeyer Award for Ideas Improving World Order, the Woodrow Wilson Award for best book published on government, politics, or international affairs, and the Robert E. Lane Award for best book in political psychology.

Personal webpage: <http://psychology.sas.upenn.edu/node/20543>

Curriculum vitae: <https://mgmt.wharton.upenn.edu/profile/1390/>

RELATED ESSAYS

Models of Revealed Preference (*Economics*), Abi Adams and Ian Crawford
Gender Segregation in Higher Education (*Sociology*), Alexandra Hendley and Maria Charles

Controlling the Influence of Stereotypes on One's Thoughts (*Psychology*), Patrick S. Forscher and Patricia G. Devine

Gender and Work (*Sociology*), Christine L. Williams and Megan Tobias Neely

The Development of Social Trust (*Psychology*), Vikram K. Jaswal and Marissa B. Drell

Genetic Foundations of Attitude Formation (*Political Science*), Christian Kandler *et al.*

Cultural Neuroscience: Connecting Culture, Brain, and Genes (*Psychology*), Shinobu Kitayama and Sarah Huff

Attitude: Construction versus Disposition (*Psychology*), Charles G. Lord

Implicit Memory (*Psychology*), Dawn M. McBride

Gender Inequality in Educational Attainment (*Sociology*), Anne McDaniel and Claudia Buchmann

Culture as Situated Cognition (*Psychology*), Daphna Oyserman

Cognitive Bias Modification in Mental (*Psychology*), Meg M. Reuland *et al.*

Born This Way: Thinking Sociologically about Essentialism (*Sociology*), Kristen Schilt

Stereotype Threat (*Psychology*), Toni Schmader and William M. Hall