
1 Introduction

Bernd Wittenbrink
Norbert Schwarz

Over the last two decades, few developments in social psychology have generated as much attention and excitement as the development of new implicit measures of attitudes, which promise to assess attitudes that respondents may not be willing to report directly or may not even be aware of themselves. The interest in these new measures has spurred significant research activity that has produced a growing number of available measures and a flurry of empirical studies concerning their effectiveness and potential limitations. This book offers a detailed introduction to this literature. Specifically, the contributions to Part I of this book describe implicit measurement procedures that have been most influential thus far. The chapters in this part outline the measures' underlying theoretical rationales, provide advice on the implementation of these measures, and review what has been learned through their use. The contributions to Part II offer diverging perspectives on implicit measures of attitudes, identify current theoretical controversies, and highlight avenues for future research. This introductory chapter provides an initial orientation for readers new to the area and offers a short preview of what is to come.

EXPLICIT AND IMPLICIT MEASURES OF ATTITUDES

Throughout the social and behavioral sciences, the dominant method of attitude measurement is the collection of explicit self-reports:

When we want to know people's attitudes toward a person, group, political issue, or product, we ask them to report them, usually by marking a rating scale or by selecting one of several response alternatives. But since the early days of attitude research, researchers have often been concerned that respondents may (sometimes) be unwilling or unable to report on their attitudes in an unbiased manner (for a review, see DeMaio, 1984). Moreover, the answers that research participants provide are highly context dependent and vary as a function of who asks, how they ask, and related variables (for reviews, see Schwarz, Groves, & Schuman, 1998; Sudman, Bradburn, & Schwarz, 1996; Tourangeau, Rips, & Rasinski, 2000). Some of the emerging context effects reflect strategic responding, whereas others reflect the cognitive and communicative processes involved in question comprehension and judgment formation. These different concerns gave rise to various methodological and theoretical answers.

Strategic Responding

Researchers' concern that participants may be unwilling to accurately report their attitudes prompted the development of three broad classes of methodological responses. One response maintains explicit self-reports as the key attitude measure and addresses respondents' presumed unwillingness to accurately report their attitudes by minimizing the incentives for socially desirable self-presentation. Relevant procedures range from simple assurances of anonymity and confidentiality to complex randomized response techniques (Bradburn, Sudman, & Warnsink, 2004). In the latter case, respondents are presented with two different questions, an innocuous one and a socially sensitive one, and draw a card that determines which one they are to answer. Given properly worded response alternatives, the interviewer remains unaware as to which question the answer pertains, thus ensuring the highest possible level of confidentiality. Other procedures create conditions that present a disincentive for socially desirable responding. For example, Sigall and Page's (1971) "bogus pipeline" technique involves convincing participants that the researcher can discern their true attitude independent of what they say, thus making lying an embarrassment. Empirically, these various techniques have been found to increase the frequency of socially undesirable answers. For example, people are more likely to admit that they enjoy pornography under randomized response conditions (Himmelfarb & Lickteig, 1982), and White participants are more likely to report that they dislike African Americans under bogus pipeline conditions (e.g., Allen, 1975). All of these procedures presume, however,

that respondents know their attitude and merely hesitate to report it. Yet deception and self-presentation may not only be directed toward others, but may also be directed toward the self (e.g., Paulhus, 1984). Perhaps people sometimes hold attitudes of which they are not aware or which they do not even want to admit to themselves.

The second class of methodological responses addresses this concern by replacing explicit self-reports of attitudes with indirect measures. Because research participants are presumably unaware of the relationship between these measures and their attitudes, indirect measures also minimize the incentives and opportunities for strategic responding. Theoretically, the use of indirect measures is based on the assumption that attitudes exert a systematic influence on people's performance on a variety of tasks and that the size of this influence can serve as an index of the underlying attitude. Not surprisingly, the theoretical assumptions made have varied widely over the history of attitude research. From the early use of projective tests (e.g., Proshansky, 1943) to the current use of response latency measures (Lane, Banaji, Nosek, & Greenwald, Chapter 3, this volume; Wittenbrink, Chapter 2, this volume) and low-tech alternatives informed by the processing assumptions of social cognition research (Vargas, Sekaquaptewa, & von Hippel, Chapter 4, this volume), the selection of indirect measures mirrors historical shifts in the underlying conceptualization of attitudes, as Vargas and colleagues note. Throughout, the usefulness of indirect measures depends on the accuracy of the bridging assumptions that link the observed response to the presumed underlying attitude, as the theoretical controversies surrounding the currently dominant response latency measures illustrate (Chapters 7–11).

Finally, a third class of methodological approaches attempts to assess research participants' evaluative responses in ways that bypass any opportunity for strategic responding, relying on the assessment of physiological reactions and brain activity (Ito & Cacioppo, Chapter 5, this volume; Olsson & Phelps, Chapter 6, this volume).

Because the latter two approaches to attitude measurement do not involve explicit self-reports, the indicators they provide can be generically referred to as implicit measures of attitudes (although some authors would prefer a more restrictive definition; see De Houwer & Moors, Chapter 7, this volume).

CONTEXT DEPENDENCY

Whereas the possibility of strategic responding does not call the existence of stable and enduring attitudes into question, the observation

that attitude reports vary as a function of numerous other contextual variables casts doubt on the assumption that attitudes are stable evaluations stored in memory. The theoretical responses to these doubts have taken a number of different forms. On one hand, attitude construal models hold that attitude reports are evaluative judgments that are made up on the spot, based on the declarative and experiential information that is accessible at the time (e.g., Schwarz & Bohner, 2001); from this perspective, the psychology of attitudes is the psychology of evaluative judgment. In contrast, others assume that context effects merely reflect noise that results from the deliberate consideration of contextual information and that attitudes are best assessed in ways that limit deliberate processing (see Ferguson & Bargh, Chapter 9, this volume, for a discussion). The most influential version of this argument conceptualizes attitudes as stored object–evaluation links that are automatically activated upon exposure to the attitude object (Fazio, 1995) and relies on evaluative priming procedures (described below and in Chapter 2) to assess the strength of the object–evaluation link. As Ferguson and Bargh (Chapter 9) emphasize, attitudes conceptualized and measured in this way are often “assumed to be contextually independent . . . , to the point that an implicit attitude measure was regarded as a potential ‘bona fide pipeline’ to people’s inner attitudes” (p. 220). Hence, implicit measures of attitudes may not only provide an answer to the problem of strategic responding, but they have also been thought to limit the context dependency of attitude measurement. Next, we turn to these measures.

IMPLICIT MEASURES OF ATTITUDES: WHAT’S AVAILABLE?

Response Time Measures

The currently most widely used implicit measures of attitudes rely on response time measurement. These measures take advantage of one of two reliable observations, namely, (1) the observation that exposure to a stimulus facilitates subsequent responses to related stimuli or (2) the observation that a stimulus is responded to more slowly when it contains multiple features that each imply a different response.

Sequential Priming Procedures

As a large body of research in cognitive psychology indicates (for a review, see Neely, 1991), exposure to a concept (e.g., *doctor*) facili-

tates the subsequent recognition of related concepts (e.g., *nurse*). A common explanation for this phenomenon holds that exposure to the initial concept (the prime) activates semantically related concepts in memory, thus reducing the time needed for their identification.

Concept priming procedures take advantage of this facilitation effect to assess a person's associations with an attitude object. For example, Wittenbrink, Judd, and Park (1997) exposed participants to African American or White primes and assessed how quickly they could identify subsequently presented trait terms of positive versus negative valence, some of which were part of the cultural stereotype about the group and some were not. The observed facilitation patterns provide information bearing on three questions: (1) Does the exposure to the group activate associated stereotypical traits, independent of their valence? If so, stereotypical traits will be recognized faster than stereotype-unrelated traits. (2) Is the automatic activation evaluatively biased; for example, are negative stereotypical traits identified more quickly than positive ones? (3) Does exposure to the group prime activate general evaluative associations, independent of their stereotypicality?

Whereas concept priming procedures present target words with descriptive meaning and use decision tasks that require participants to identify the word, *evaluative priming* procedures (Fazio, Sanbonmatsu, Powell, & Kardes, 1986) present target words with general evaluative meanings (e.g., awful, pleasant) and ask participants to judge the words' evaluative connotation (good or bad). Of interest is whether exposure to an attitude object facilitates the evaluative response to negative or positive target words. Thus, evaluative priming assesses whether an attitude object triggers an automatic evaluation, whereas concept priming assesses descriptive associations that may have evaluative content.

Wittenbrink (Chapter 2, this volume) reviews these procedures, provides advice on their implementation, and summarizes representative findings.

Response Competition Procedures

Whereas the preceding procedures take advantage of priming effects, a second class of response time procedures is based on interference effects that may occur when different features of an attitude object imply different responses. The best known of these procedures is the Implicit Association Test (IAT) (Greenwald, McGhee, & Schwartz, 1998), reviewed by Lane, Banaji, Nosek, and Greenwald (Chapter 3, this volume).

The IAT presents two discrimination tasks that are combined in specific ways across a sequence of five steps. To assess attitudes toward African Americans and European Americans, for example, the first discrimination task may present names that are typical for the respective group and then ask participants to categorize each name as “White” versus “Black.” They do so by pressing a response key assigned to “White” with the left hand or a response key assigned to “Black” with the right hand. Next, the second discrimination task presents words with pleasant (e.g., love) or unpleasant (e.g., poison) connotations, which participants classify as positive versus negative by pressing the left or right response key. At the third step, these two tasks are superimposed and participants press the left key when either a White name or a pleasant word is shown, but the right key when either a Black name or an unpleasant word is shown. As in the above facilitation paradigms, this task is easier when evaluatively associated categories share the same response key—for example, when White participants press the left key to categorize White names and pleasant words. Going beyond this assessment of response facilitation, the IAT involves two more steps. At the fourth step, the assignment of keys to White and Black names is reversed, so that participants who first used the left key for White names now use the left key for Black names. Finally, the two discrimination tasks are again superimposed, resulting in an assignment of “Black” and “pleasant” to the left response key and “White” and “unpleasant” to the right response key.

Of interest is the speed with which participants can perform the two superimposed discrimination tasks at step 3 and step 5. Do participants respond faster when a given response key pertains either to the pairing of White names + pleasant words or Black names + unpleasant words (step 3) than when this pairing is reversed and a given response key pertains either to White names + unpleasant words or Black names + pleasant words (step 5)? In the present example, a faster response at step 3 than at step 5 is thought to indicate that White names and positive evaluations, and Black names and negative evaluations, are more strongly associated than the reverse pairings.

Lane and colleagues (Chapter 3, this volume) review the underlying logic, report representative findings, and provide hands-on advice for the implementation and scoring of the IAT. Related response competition tasks include the Go/No-go Association Task (GNAT; Nosek & Banaji, 2001) and the Extrinsic Affective Simon Task (EAST; De Houwer, 2003).

Paper-and-Pencil Measures

Whereas the preceding measures require a high degree of instrumentation and technical sophistication, other implicit measures of attitudes are decidedly low-tech. Vargas, Sekaquaptewa, and von Hippel (Chapter 4, this volume) provide an informative review of a wide range of such low-tech measures and place them in the context of the history of attitude research.

Drawing on insights from social cognition research, some of these measures take advantage of the observation that attitudes and expectations have systematic effects on individuals' information processing. For example, people are more likely to spontaneously explain events that disconfirm rather than confirm their expectations (e.g., Hastie, 1984), suggesting that the amount of explanatory activity can serve as an indirect measure of a person's expectations. The Stereotypic Explanatory Bias (SEB) measure developed by Sekaquaptewa, Espinoza, Thompson, Vargas, and von Hippel (2003) builds on this observation and uses the number of explanations generated in response to stereotype-consistent versus stereotype-inconsistent behaviors as an implicit measure of stereotyping. Similarly, people describe expected or stereotype-consistent behaviors in more abstract terms than unexpected or stereotype-inconsistent behaviors, a phenomenon known as the Linguistic Intergroup Bias (LIB; see, e.g., Maass, Salvi, Arcuri, & Semin, 1989). The size of this bias can again be used as an indirect measure to gauge the underlying expectations. As Vargas and colleagues (Chapter 4, this volume) review, such measures have been found to predict prejudiced behaviors, although little is known about their psychometric qualities. Given their ease of use, the various measures reviewed by Vargas and colleagues deserve more systematic methodological exploration.

Physiological Responses and Brain Activity

Because of their involuntary and hard-to-control nature, physiological correlates of evaluative responses have long been of interest to attitude researchers who doubted respondents' explicit self-reports. Whereas early attempts to rely on the galvanic skin response (e.g., Rankin & Campbell, 1955) or on pupillary dilation or restriction (e.g., Hess, 1965) as indirect measures met with little success, recent progress in social psychophysiology and neuroscience suggests a more optimistic outlook, although a one-to-one mapping of neural and psychological processes is unlikely, as Ito and Cacioppo (Chapter 5, this volume) emphasize.

Ito and Cacioppo (Chapter 5, this volume) provide a tutorial overview of the available measures, ranging from autonomic responses, like cardiovascular and electrodermal activity, to facial electromyography and startle eyeblink modification or measures of brain activity, like functional magnetic resonance imaging and event-related brain potentials. Their review highlights the numerous intricacies of these measures, which require considerable specialized expertise for their implementation and interpretation. Olsson and Phelps's (Chapter 6, this volume) discussion of what we can and cannot learn from neuroimaging complements Ito and Cacioppo's overview and summarizes the neural underpinnings of social evaluations.

Summary

In combination, this first set of contributions (Chapters 2–6) provides an overview of the current state of the art in the implicit measurement of attitudes. These chapters review the currently available measures, offer advice on their implementation and interpretation, and summarize representative research findings. The remaining chapters provide different theoretical perspectives on the operation of these measures and address current controversies.

PERSPECTIVES AND CONTROVERSIES

What Makes a Measure “Implicit”?

Above, we distinguished between explicit and implicit measures of attitudes by virtue of their transparency and potential for strategic responding. In Chapter 7, this volume, De Houwer and Moors provide a thoughtful discussion of what qualifies a measure as “implicit.” They adopt a more restrictive conceptualization that defines implicit measures as “measurement outcomes that reflect the to-be-measured construct by virtue of processes that are uncontrolled, unintentional, goal independent, purely stimulus driven, autonomous, unconscious, efficient, or fast” (pp. 188–189). Although all of the measures reviewed in Chapters 2–6 meet some of these criteria, few are likely to meet all of them. In fact, the extent to which even the most widely used measures meet some of these criteria is currently unknown, and De Houwer and Moors outline a research program that addresses these issues. Gawronski and Bodenhausen's (Chapter 11, this volume) discussion of conceptual and terminological ambiguities echoes these concerns.

Reiterating De Houwer and Moors's (Chapter 7, this volume) call for a more detailed analysis of the processes underlying implicit measures of attitudes, Wentura and Rothermund (Chapter 8, this volume) offer a plea for more basic research on the experimental paradigms on which implicit attitude measures are based. Their chapter raises several important questions about how response latency measures actually work and what exactly they measure. In general, it is assumed that latency-based measures capture associations between constructs and evaluations. In the case of the IAT, for example, a person is thought to respond faster to one set of key pairings because he or she holds relatively stronger associations for this pair than for the other. Wentura and Rothermund note that not all findings are perfectly consistent with such an associative account and, instead, offer an alternative explanation. In particular, they suggest that IAT effects reflect differences in the relative salience of the employed response categories. Although such salience differences themselves may be influenced by evaluative associations, there are other factors that likely impact salience as well and thus might limit the validity of the IAT as a measure of associative strength between an attitude object and an evaluation. Hence, the chapter by Wentura and Rothermund highlights the importance of future research on the mechanisms underlying implicit measurement procedures. Clarifying the exact mechanism by which these measures operate should improve our understanding of how to use and design effective implicit measurement tools.

Are Implicit Measures Context Dependent?

As Ferguson and Bargh (Chapter 9, this volume) highlight, the initial hope that responses to implicit measures of attitudes that limit deliberation may be less context dependent than responses to explicit attitude questions has not been supported. Instead, evaluative and conceptual priming procedures (Wittenbrink, Chapter 2, this volume) and response competition procedures like the IAT (Lane et al., Chapter 3, this volume) show pronounced context effects that usually parallel the context effects observed on explicit attitude measures. For example, Dasgupta and Greenwald (2001) observed that exposure to pictures of liked African Americans and disliked European Americans resulted in shifts on a subsequent IAT that parallel the effects of exposure to liked or disliked exemplars on explicit measures of attitudes (e.g., Bodenhausen, Schwarz, Bless, & Wänke, 1995). Similarly, Wittenbrink, Judd, and Park (2001) found that the same Black face primes elicited more negative automatic responses when the faces were presented on

the background of an urban street scene rather than a church scene. Blair (2002) provides an extensive review of related findings.

Of particular interest is Lowery, Hardin, and Sinclair's (2001) observation that participants provided less negative automatic evaluations of African Americans when the experimenter was Black rather than White. This finding parallels the observed influence of interviewer race and ethnicity in the survey research literature (e.g., Hatchett & Schuman, 1976; Weeks & Moore, 1981). Yet the low transparency of Lowery and colleagues' implicit attitude measures makes it unlikely that participants' favorable attitude toward African Americans was based on deliberate strategic responding. Instead, it raises the possibility that experimenters and interviewers serve as highly accessible positive exemplars, thus increasing the favorability of the response as observed in studies that used names or pictures as exemplar primes (e.g., Bodenhausen et al., 1995; Dasgupta & Greenwald, 2001). This possibility would have far-reaching implications for the conceptualization of race of interviewer effects and the assumed role of socially desirable responding in their emergence.

To account for the context dependency of implicit measures, Ferguson and Bargh (Chapter 9, this volume) suggest that automatic attitudes are responses to object-centered contexts rather than to the attitude object in isolation. They conceptualize the underlying process in terms of a connectionist system, consistent with Smith and Conrey's (Chapter 10, this volume) assertion that "mental representations are states and not things."

From Smith and Conrey's (Chapter 10) perspective, the context sensitivity of implicit as well as explicit measures reflects the dynamic and context-sensitive nature of the mental representations on which evaluative responses are based. Only context-sensitive representations allow "the mind to respond efficiently and accurately to a constantly changing environment that calls for situated knowledge and behaviors" (p. 256). Hence, implicit measures of attitudes do not (solely) reflect previously learned object-evaluation links, but capture the current evaluative response to the attitude object in its present context, consistent with the observation that automatic evaluations can be obtained for novel objects, for which no previously acquired object-attitude links are stored in memory (Duckworth, Bargh, Garcia, & Chaiken, 2002).

Where To?

In the concluding Chapter 11, this volume, Gawronski and Bodenhausen identify important open issues, ranging from diverging theoretical

conceptualizations of attitudes, to the (often lacking) correspondence between these conceptualizations and measurement procedures, and the conditions under which results obtained with different implicit measures may or may not converge. Their discussion provides welcome conceptual clarifications and suggests an ambitious but promising agenda for future research. In particular, Gawronski and Bodenhausen advocate the integration of existing small-scale explanations for how implicit assessment procedures work with broader, large-scale theories on the determinants of judgment and behavior and the interplay of affect and cognition. At present, accounts of the procedural underpinnings of implicit measures tend to involve concrete operational constructs that do not readily translate into the more abstract terminology of broader theories of social judgment and behavior. An integration of these divergent perspectives could prove immensely useful for furthering our understanding of the way attitudes determine behavior, as well as improve our ability to assess these influences.

Empirical and theoretical research in this area is progressing at a rapid pace, and an increasing number of researchers make use of implicit measures of attitudes. We hope that this book's introductory tutorials will provide useful guidance on the use and implementation of these measures and that the accompanying critical perspectives foster awareness of the numerous remaining ambiguities and open issues.

REFERENCES

- Allen, B. P. (1975). Social distance and admiration reactions of "unprejudiced" whites. *Journal of Personality*, *43*, 709–726.
- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, *6*, 242–261.
- Bodenhausen, G. V., Schwarz, N., Bless, H., & Wänke, M. (1995). Effects of atypical exemplars on racial beliefs: Enlightened racism or generalized appraisals? *Journal of Experimental Social Psychology*, *31*, 48–63.
- Bradburn, N., Sudman, S., & Wansink, B. (2004). *Asking questions* (2nd ed.). San Francisco: Jossey-Bass.
- Dasgupta, N., & Greenwald, A. G. (2001). On the malleability of automatic attitudes: Combating automatic prejudice with images of liked and disliked individuals. *Journal of Personality and Social Psychology*, *81*, 800–814.
- De Houwer, J. (2003). The Extrinsic Affective Simon Task. *Experimental Psychology*, *50*, 77–85.
- DeMaio, T. J. (1984). Social desirability and survey measurement: A review. In C. F. Turner & E. Martin (Eds.), *Surveying subjective phenomena* (Vol. 2, pp. 257–281). New York: Russell Sage.

- Duckworth, K. L., Bargh, J. A., Garcia, M., & Chaiken, S. (2002). The automatic evaluation of novel stimuli. *Psychological Science*, *13*, 513–519.
- 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., 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.
- Greenwald, A. G., McGhee, D. E., & Schwarz, J. L. K. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology*, *74*, 1464–1480.
- Hastie, R. (1984). Causes and effects of causal attribution. *Journal of Personality and Social Psychology*, *46*, 44–56.
- Hatchett, S., & Schuman, H. (1976). White respondents and race-of-interviewer effects. *Public Opinion Quarterly*, *39*, 523–528.
- Hess, E. H. (1965). Attitude and pupil size. *Scientific American*, *212*, 46–54.
- Himmelfarb, S., & Lickteig, C. (1982). Social desirability and the randomized response technique. *Journal of Personality and Social Psychology*, *43*, 710–717.
- Lowery, B. S., Hardin, C. D., & Sinclair, S. (2001). Social influence on automatic racial prejudice. *Journal of Personality and Social Psychology*, *81*, 842–855.
- Maass, A., Salvi, D., Arcuri, L., & Semin, G. (1989). Language use in intergroup contexts: The linguistic intergroup bias. *Journal of Personality and Social Psychology*, *57*, 981–993.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading: Visual word recognition* (pp. 264–336). Hillsdale, NJ: Erlbaum.
- Nosek, B. A., & Banaji, M. R. (2001). The Go/No-go Association Task. *Social Cognition*, *19*(6), 625–666.
- Paulhus, D. L. (1984). Two-component models of socially desirable responding. *Journal of Personality and Social Psychology*, *46*(3), 598–609.
- Proshansky, H. M. (1943). A projective method for the study of attitudes. *Journal of Abnormal and Social Psychology*, *38*, 393–395.
- Rankin, R. E., & Campbell, D. T. (1955). Galvanic skin response to negro and white experimenters. *Journal of Abnormal and Social Psychology*, *51*, 30–33.
- Schwarz, N., & Bohner, G. (2001). The construction of attitudes. In A. Tesser & N. Schwarz (Eds.), *Blackwell handbook of social psychology: Intraindividual processes* (pp. 436–457). Malden, MA: Blackwell.
- Schwarz, N., Groves, R., & Schuman, H. (1998). Survey methods. In D. Gilbert, S. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (4th ed., Vol. 1, pp. 143–179). New York: McGraw-Hill.
- Sekaquaptewa, D., Espinoza, P., Thompson, M., Vargas, P., & von Hippel, W. (2003). Stereotypic explanatory bias: Implicit stereotyping as a predictor of discrimination. *Journal of Experimental Social Psychology*, *39*, 75–82.
- Signall, H., & Page, R. (1971). Current stereotypes: A little fading, a little faking. *Journal of Personality and Social Psychology*, *18*, 247–255.

- Sudman, S., Bradburn, N. M., & Schwarz, N. (1996). *Thinking about answers: The application of cognitive processes to survey methodology*. San Francisco: Jossey-Bass.
- Courageous, R., Rips, L. J., & Resins, K. (2000). *The psychology of survey response*. Cambridge, UK: Cambridge University Press.
- Weeks, M. F., & Moore, R. P. (1981). Ethnicity-of-interviewer effects on ethnic respondents. *Public Opinion Quarterly*, 45, 245–249.
- Writtenbrink, 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.
- Writtenbrink, B., Judd, C. M., & Park, B. (2001). Spontaneous prejudice in context: Variability in automatically activated attitudes. *Journal of Personality and Social Psychology*, 81, 815–827.