

Research Article

Implicit Attitude Generalization Occurs Immediately; Explicit Attitude Generalization Takes Time

Kate A. Ranganath and Brian A. Nosek

University of Virginia

ABSTRACT—*People are able to explicitly resist using knowledge about one person to evaluate another person from the same group. After learning about positive and negative behaviors performed by one individual from each of two different groups, participants were introduced briefly to new individuals from the groups. Implicit evaluations of the original individuals readily generalized to the new individuals; explicitly, participants resisted such generalization. Days later, both implicit and explicit evaluations of the original individuals generalized to the new individuals. The results suggest that associative links (e.g., shared group membership) are sufficient for implicit attitude generalization, but deliberative logic (e.g., individual group members are not necessarily the same) can reduce explicit generalization by association. When knowledge distinguishing who did what is unavailable, such as after forgetting, associative knowledge provides the basis of explicit evaluation. We conclude that a simple association linking one individual to another can produce implicit attitude generalization immediately and explicit attitude generalization eventually.*

Moral intuition and legal theory indicate that it is unacceptable to judge a person on the basis of the actions of another. Ancient and modern jurisprudence reject the notion of guilt by association (see Banaji & Bhaskar, 2000). Individuals are not deliberately punished for the crimes of others who share the same group membership—whether the group is defined by kinship,

friendship, or social category. Nevertheless, everyday thinking and judging provide occasions for evaluative generalization in which the actions of one individual are used in evaluating another.

Crawford, Sherman, and Hamilton (2002) demonstrated that transferring the traits perceived about one individual to another group member depends on perceived *group entitativity*—the extent to which a collection of individuals is perceived as being a unified entity (Campbell, 1958). Participants were given positive or negative information about an individual group member and were then given an opportunity to apply that information to other group members (or not). The key manipulation was whether the group to which the individuals belonged was described as having high entitativity (similar people with shared background, attitudes, and personalities) or low entitativity (diverse people with different backgrounds, attitudes, and personalities). When the group was perceived as being highly entitative, participants transferred traits from one group member to another. But when the social group was thought to be large and diverse, participants recognized that one individual is not representative of all group members and resisted generalizing their attitude from one person to another person in the same group.

A nuanced understanding of group diversity overrides a more basic relation—that individuals are associated by simply belonging to the same group. To avoid using evaluations of one person in judging another requires a deliberate judgment that group association is not a sufficient basis for generalization. Despite the conscious thought involved in resisting attitude generalization, the mental presence of a group association might be sufficient for observers to generalize even when they have no intention of doing so. In other words, even when attitude generalization is explicitly resisted, associations between group members may still exist and lead to generalization anyway.

Address correspondence to Kate Ranganath, Department of Psychology, University of Virginia, PO Box 400400, Charlottesville, VA 22904, e-mail: ranganath@virginia.edu.

THE SUBTLE INFLUENCE OF ASSOCIATIVE RELATIONS

A variety of modern dual-process models suggest that qualitatively distinct cognitive processes are involved in social evaluation. Gawronski and Bodenhausen (2006) and other researchers (Sloman, 1996; Smith & DeCoster, 2000; Strack & Deutsch, 2004) have suggested that both associative and propositional processes contribute to evaluation. Associative evaluations are activated automatically when a person encounters a relevant stimulus. Propositional processes, in contrast, influence deliberate, reflective decisions that involve the manipulation of symbolically represented rules. Gawronski and Bodenhausen have further proposed that the most important distinction between associative and propositional reasoning is that associative relations can be activated regardless of whether a person believes those associations are accurate or inaccurate, whereas propositional processes allow for the deliberation involved in conscious declarations of truth or falsity.

For example, according to classical and modern conceptions of cognitive balance, if one learns that Oscar has many negative characteristics and then learns that Elmo belongs to the same group that Oscar belongs to, Elmo should become associated with negativity because of his association with Oscar (Greenwald et al., 2002; Heider, 1958). However, consciously applying propositional logic that Oscar's actions should not reflect on Elmo should inhibit associative transference. Deliberation allows for the rejection of the associative generalization and for Elmo to be judged on his own merits. Recent work showing the distinctiveness of implicit (associative) and explicit (deliberative) evaluations provides evidence that people can possess multiple evaluations of a single target—even evaluations they would reject as false if given the opportunity to deliberate (Greenwald & Banaji, 1995; Nosek, 2007).

In the present study, participants read about behaviors performed by two individuals who belonged to distinct, heterogeneous groups. One individual performed mostly negative behaviors, and the other performed mostly positive behaviors. Immediately after reading about these two individuals, participants were introduced briefly to two new individuals from the same groups. We predicted that the induced evaluations of the original individuals would generalize to the new individuals associatively, as measured by the Implicit Association Test (IAT; Greenwald, McGhee, & Schwartz, 1998), but that participants would resist generalizing from the original to the new individuals when given the opportunity to deliberate via self-reported evaluations.

RESOURCES PREVENTING GENERALIZATION FADE WITH TIME

The existing studies showing that people resist generalizing attitudes from one person to another have typically measured

evaluations immediately after an attitude induction. Observers in these circumstances have ample resources with which to deliberate and prevent generalization. In particular, clear memory for who did what is necessary for observers to keep their evaluations straight and enables them to prevent the actions of one individual from influencing their judgments of another.

With time, however, the details that maintain clear distinctions between the behaviors of different individuals are forgotten (Johnson, Hashtroudi, & Lindsay, 1993). Brainerd and Reyna's fuzzy-trace theory, for example, distinguishes between gist memory (memory for the general meaning of informational inputs) and verbatim memory (more exact representation of the specific content of informational inputs; Reyna & Brainerd, 1995). According to this theory, verbatim memory is forgotten and becomes inaccessible quickly, leaving people to rely on gist memory (Clark & Clark, 1977).

The quicker loss of verbatim memory relative to gist memory suggests that, over time, the specific information used to deliberately distinguish group members will fade. If the associations between group members still exist in memory without the conscious individuating information, then the passage of time should increase the difficulty of avoiding associative generalization from one person to another. To test this possibility, we reassessed implicit and explicit evaluations toward the novel individuals days after the induction. We expected that participants would be able to resist explicit (but not implicit) generalization when tested immediately, but that they would not be able to explicitly resist generalization following a delay. Thus, we predicted that after the delay, both implicit and explicit evaluations of the original individuals would generalize to the new individuals.

STUDY OVERVIEW

The study took place in two sessions. In the first session, participants read about behaviors performed by two individuals, Reemolap and Vabbenif (adapted from Gregg, 2000), each of whom was described as belonging to a large and diverse social group (low in entitativity). One of the individuals performed primarily positive behaviors, and the other performed primarily negative behaviors. Next, participants read brief introductions to two new individuals, Bosaalap and Ibbonif, who belonged to the same social groups as the original individuals. The information about the new individuals was minimal, and pretesting showed it to be relatively neutral. We then used the IAT and semantic differential ratings to measure participants' implicit and explicit attitudes, respectively, toward either the original individuals or the new individuals. We manipulated the individuals to be evaluated between participants so that the act of reporting attitudes toward one pair would not influence reports of attitudes toward the other. The IAT is widely used to measure conditioned associations (Baccus, Baldwin, & Packer, 2004; Gregg, Seibt, & Banaji, 2006; Olson & Fazio, 2001; Rydell, McConnell, Mackie, & Strain,

2006), and, because of its demonstrated reliability and sensitivity, it was particularly useful for assessing associations in this study (for a review of the psychometric properties of the IAT, see Nosek, Greenwald, & Banaji, 2007).

We predicted that people would explicitly refrain from using the information about the original group members in making immediate judgments of the new group members; this finding would replicate demonstrations that people can recognize that one individual's behavior is not representative of all group members if their group is large and diverse (Crawford et al., 2002; McConnell, Sherman, & Hamilton, 1997). In contrast, we predicted that implicit attitudes toward the new individuals would be as strong as if the new individuals were the targets of the original induction. That is, we expected the induced evaluations of the original individuals to immediately transfer associatively to the new individuals as a consequence of their shared group membership.

Three days after the first study session, we invited participants to return to complete the implicit and explicit attitude measures again. We hypothesized that days after the initial learning episode, participants would show not only the implicit attitude generalization predicted for the first session, but also full explicit attitude generalization. That is, after the delay, we expected the attitudes induced toward the original individuals to provide the basis for explicitly evaluating the new individuals.

METHOD

Participants

Participants were volunteers at the Project Implicit research Web site (<https://implicit.harvard.edu>). They were randomly assigned to this study, which was part of an available pool of approximately 100 studies. At the end of the first session, participants were asked if they were willing to be contacted again for a follow-up session. Of the 831 people who participated in both sessions, 684 (82%) completed all of the study materials.¹ For more information about this virtual laboratory, see Nosek (2005) and Nosek et al. (2007).

¹Participants who completed the second session were slightly older ($M = 36.15$ years, $SD = 12.86$) than those who did not ($M = 32.54$ years, $SD = 12.56$), $t(1681) = 4.88$, $p_{\text{rep}} = .98$, $d = 0.12$. Also, the percentage of women was significantly higher in the group who completed the second session (71%) than in the group who did not complete that session (62%), $t(1679) = 4.26$, $p_{\text{rep}} = .98$, $d = 0.11$. On a scale ranging from -3 (*very liberal*) to 3 (*very conservative*), those participants who completed the second session were significantly more liberal ($M = -1.17$, $SD = 1.62$) than those who did not ($M = -0.67$, $SD = 1.72$). However, there were no significant differences in substantive effects in the first session when we compared participants who did and who did not complete the second session: Implicit attitudes did not differ significantly, $F(1, 1654) = 1.80$, $p_{\text{rep}} = .74$, $d = 0.04$, and neither did explicit attitudes, $F(1, 1652) = 0.02$, $p_{\text{rep}} = .20$, $d = 0.0005$. Because results in the first session did not differ between the two groups, and because one of our primary hypotheses concerned differences in immediate and delayed attitude transfer, in this article we present results for those participants who completed both sessions.

Measures

IAT (Greenwald et al., 1998)

The IAT assesses associations among two concept categories (e.g., Reemolap and Vabbenif) and two evaluative attributes (e.g., good and bad, the attributes used in this study) by requiring that participants categorize stimulus items representing the four categories as quickly as possible using two keys of a computer keyboard. Participants used the “e” and “i” keys to classify items one at a time into the corresponding superordinate categories identified on the right or left side of the screen. At each session, participants completed an IAT for either the original people or the new people.

The IATs consisted of seven blocks of trials, following the recommendation of Nosek, Greenwald, and Banaji (2005). We analyzed the response latencies using the D algorithm (Greenwald, Nosek, & Banaji, 2003). Response latency was calculated from the beginning of the trial until the time of a correct response, and latencies less than 400 ms were removed. A positive D score indicates an implicit preference consistent with the attitude induction—that is, a preference for the individual who was associated with positive information (original-people condition) or a preference for the new individual who belonged to the same group as the original person associated with positive information (new-people condition).

Self-Report Measures

At each session, explicit attitudes were measured with 8-point semantic differential ratings for either the original individuals or the new individuals. The following traits were rated: unintelligent-intelligent, unpopular-popular, unlikeable-likeable, unfriendly-friendly, and unpleasant-pleasant. Participants also reported their global liking for each of the individuals (“How much do you like . . . ?”) on a 7-point scale ranging from *dislike very much* to *like very much*. Reliability among ratings was high (Cronbach's $\alpha = .89$), so they were averaged into a single rating for each individual.

Manipulation Check

At the end of the first session, participants were asked to identify the four individuals to whom they were introduced by matching names and pictures. Two percent of participants failed this manipulation check. The substantive results were unchanged when these participants were excluded from analysis.

Procedure

First Session

The attitude-induction procedure was designed to create a clear preference for one individual over the other. Participants viewed a series of sentences describing behaviors performed by Reemolap and Vabbenif. Each sentence was accompanied by a

picture of the individual, counterbalanced across participants, who performed that behavior. Reemolap performed predominantly positive behaviors (12 positive, 4 negative), and Vabbenif performed predominantly negative behaviors (12 negative, 4 positive). The name-behavior pairs were adapted from Hamilton and Gifford (1976; e.g., Reemolap visited a sick friend in the hospital, and Vabbenif parked in a space reserved for the handicapped). Each name-behavior pair appeared on a computer screen for 3 s. The induction procedure took approximately 3 min.

Next, participants were introduced to the new people. Bosaalap and Ibbonif were described as belonging to the same groups as Reemolap and Vabbenif—Laapians and Niffians, respectively. Both groups were described as diverse groups with many members who perform all different types of behaviors. The information distinguishing the new individuals consisted of a cartoon picture showing both and these descriptions:

Ibbonif is a sculptor and very much enjoys gardening, biking, and playing card games. Ibbonif is kind and thoughtful, but tends to be slightly greedy at times.

Bosaalap is a painter and very much enjoys cooking, hiking, and listening to music. Bosaalap is warm and considerate, but tends to be slightly dishonest at times.

In pretesting, these descriptions were judged to be evaluatively equivalent. The cartoon images of the original and new individuals had a general resemblance consistent with the group memberships. Picture-description pairings were randomized across participants and did not affect the results.

Finally, participants were randomly assigned to complete the implicit and explicit attitude measures about either the original person (Reemolap and Vabbenif) or the new people (Bosaalap and Ibbonif). The order of the implicit and explicit measures was randomized across participants and did not influence the results.

Second Session

At the end of the first session, 92% of participants indicated their willingness to be contacted by e-mail for a follow-up session. These participants were e-mailed a link to the follow-up session 3 days later (days between sessions: $M = 10$, $SD = 7$, $Mdn = 6$, mode = 5). The length of time between the first and second sessions did not moderate any results.

Participants were randomly assigned to complete the implicit and explicit attitude measures toward either the original people or the new people. Thus, roughly half the participants completed the attitude measures for the same people at the first and second sessions, and the other half completed the attitude measures for different people at the two sessions. Whether they completed attitude measures for the same or different people at the first and second sessions did not qualify the results.

RESULTS

Attitudes Were Induced

Self-reported attitudes toward the original people, the targets of the attitude induction, were in a direction consistent with the induction materials. The individual who performed primarily positive behaviors ($M = 4.26$, $SD = 1.03$) was explicitly preferred over the individual who performed primarily negative behaviors ($M = 3.66$, $SD = 0.97$), $t(291) = 11.46$, $p_{rep} = .99$, Cohen's $d = 0.65$. Implicit attitudes toward the original people were also consistent with the induction materials. The individual who performed primarily positive behaviors was implicitly preferred over the individual who performed primarily negative behaviors ($M = 0.07$, $SD = 0.42$), $t(291) = 3.02$, $p_{rep} = .98$, $d = 0.17$.

Explicitly, Generalization Was Resisted; Implicitly, Evaluations Generalized to New Individuals

Participants resisted using information about one individual in their immediate explicit evaluations of another individual from the same group. The attitude induction had a significantly stronger effect on self-reported attitudes toward the original people, the targets of the induction, than on self-reported attitudes toward the new people (the mean difference between the attitudes toward the person who performed primarily positive behaviors and the person who performed primarily negative behaviors was 0.60, $SD = 1.72$, for the original people and 0.36, $SD = 1.15$, for the new people), $t(566) = 5.72$, $p_{rep} = .99$, $d = 0.25$.

Notably, explicit attitude generalization was not entirely avoided. If there had been no generalization of induced attitudes, we would have observed no preference. However, attitudes toward the new individuals did show a preference consistent with what was learned about their fellow group members; that is, participants preferred the new individual who belonged to the same group as the original person who performed primarily positive behaviors ($M = 4.44$, $SD = 0.79$) over the new individual who belonged to the same group as the original person who performed primarily negative behaviors ($M = 4.09$, $SD = 0.80$), $t(277) = 5.17$, $p_{rep} = .99$, $d = 0.29$. One factor that may have prevented complete resistance to generalization was that attitudes toward the original and new individuals were measured between participants, which reduced the opportunity for participants to deliberately distinguish their attitudes toward the new people from their attitudes toward the original people. Nonetheless, participants explicitly resisted generalization; the effect of the attitude induction on the explicit ratings was 70% smaller for the new people than for the original people.

The induced implicit attitudes toward the original people generalized completely to the new people. The new individual who shared group membership with the original person who performed primarily positive behaviors was implicitly preferred over the new individual who shared group membership with the

original person who performed primarily negative behaviors ($M = 0.08$, $SD = 0.45$), $t(273) = 2.88$, $p_{\text{rep}} = .98$, $d = 0.18$. There was no difference between the effectiveness of the induction on attitudes toward the original people ($M = 0.07$, $SD = 0.42$) and attitudes toward the new people ($M = 0.08$, $SD = 0.45$), $t(562) = -0.11$, $p_{\text{rep}} = .18$, $d = -0.004$, even though the new individuals were not involved in the induction procedure. Participants were unable to prevent the information about the original individuals from influencing their implicit responses toward the new individuals. They automatically and immediately treated the individual group members as if they were interchangeable.

Induced Attitudes Persisted Over Time

At the second session, self-reported attitudes toward the original people, the targets of the attitude induction, continued to show a preference consistent with the induction materials. The individual who performed primarily positive behaviors ($M = 4.15$, $SD = 0.93$) was preferred over the individual who performed primarily negative behaviors ($M = 3.71$, $SD = 0.90$), $t(273) = 5.56$, $p_{\text{rep}} = .98$, $d = 0.34$. Implicit attitudes toward the original people also persisted over time. The individual who performed primarily positive behavior was implicitly preferred over the individual who performed primarily negative behaviors ($M = 0.08$, $SD = 0.46$), $t(273) = 2.82$, $p_{\text{rep}} = .97$, $d = 0.17$.

Implicit and Explicit Attitudes Generalized After a Delay

Recall that when explicit attitudes were measured immediately after training, participants resisted transferring their evaluative knowledge of one individual to another from the same group. At the second session, however, self-reported attitudes toward the original people generalized completely to the new people. The new person who shared group membership with the original person who performed primarily positive behaviors ($M = 4.17$, $SD = 0.93$) was preferred over the new person who shared group membership with the original person who performed primarily negative behaviors ($M = 3.72$, $SD = 0.90$), $t(288) = 6.58$, $p_{\text{rep}} = .98$, $d = 0.39$. There was no difference between the effectiveness of the induction on evaluations of the original and new individuals (the mean difference between the attitudes toward the person who performed primarily positive behaviors and the person who performed primarily negative behaviors was 0.45, $SD = 1.46$, for the original people and 0.40, $SD = 1.36$, for the new people), $t(559) = 0.24$, $p_{\text{rep}} = .28$, $d = 0.01$.

Also, the generalized implicit attitudes that were induced immediately persisted over time. At the second testing session, the implicit preference between the new people showed an effect consistent with the induction materials describing the original people ($M = 0.12$, $SD = 0.46$), $t(566) = 4.36$, $p_{\text{rep}} = .99$, $d = 0.26$. Again, there was no difference between the effectiveness of the induction on evaluations of the original ($M = 0.08$, $SD = 0.46$) and new ($M = 0.12$, $SD = 0.46$) individuals, $t(566) = -1.03$, $p_{\text{rep}} = .64$, $d = -0.04$.

Finally, the correlation between implicit and explicit attitudes was significantly stronger at the delayed test ($r = .28$) than at the immediate test ($r = .16$), $Z = 2.31$, $p_{\text{rep}} = .96$. This finding is consistent with our prediction that implicit associations would have a stronger influence on explicit judgment as the details of the induction faded in memory.

DISCUSSION

Even when people deliberately resist using the behaviors of one individual to make explicit judgments about another individual from the same group, the simple association due to shared group membership is sufficient for implicit evaluations of one group member to generalize to another. Further, generalization occurs even explicitly when evaluations are measured days after the original induction. Despite participants' ability to consciously resist using group information as a basis of evaluation immediately, associations between group members still exist in memory and influence evaluation as time passes.

Balance principles predict that if one person who is viewed negatively is associated with a second person by group membership, then the second person is likely to inherit the negative evaluation (Greenwald et al., 2002; Heider, 1958). We assume that resisting transferring attitudes from one person to another requires the ability to differentiate distinct knowledge about group members and the cognitive capacity to deliberately refrain from generalization. As information individuating two group members fades, resisting generalization becomes more difficult. As suggested by fuzzy-trace theory (Reyna & Brainerd, 1995), the details that distinguish knowledge about one group member from knowledge about another are forgotten, leaving conscious evaluation reliant only on the associations that are available in memory.

Like the other studies that have used the IAT to measure conditioned associations, ours rests on the assumption that the IAT measures associations (see also Greenwald, Nosek, Banaji, & Klauer, 2005). Rothermund and Wentura (2004) argued that the IAT may be influenced by salience asymmetries as well. However, it is unlikely that our conditioning paradigm induced salience asymmetries because we used novel groups and presented the same amount of information about each individual. The only difference between the individuals was that one was associated with positive behaviors and the other with negative behaviors.

The present results inspire a host of new questions concerning boundary conditions and variables that may moderate the influence of associative generalization on evaluation. For example, explicit generalization across individuals within a group should increase as a function of perceived similarity, but does implicit generalization occur even following trivial means of associating individuals? Also, this paradigm might be adapted to investigate when individuating or generalizing processes will dominate, and under what conditions automatic processes will lead to generalization despite conscious intentions otherwise.

The present work suggests that associative processes influence propositional processes in social evaluation and is a significant step in specifying the importance of associative relationships.

Even if people believe that it is unacceptable to judge people on the basis of the actions of others, associative relations provide a subtle means of inducing guilt or accolade by association. This research demonstrated immediate and automatic evaluation of little-known social targets based on evaluations of associated individuals. An encounter with one member of a group can have lasting effects on evaluations of other members of the same group. In social judgment, generalizing evaluations from one person to others provides a simple heuristic for assessment. However, the potential cost is that individuating information that would enable more accurate assessment may be missed because of a propensity to generalize qualities that group members may not hold in common.

Acknowledgments—We would like to thank Tony Greenwald, Lee Hendricks, Colin Tucker Smith, and the Implicit Cognition Lab at the University of Virginia for their helpful comments. We would also like to thank Jeff Hansen for his technical expertise and David Hamilton for his generosity in sharing study materials. This research was supported by a grant from the National Institute of Mental Health (R01 MH68447) to Brian Nosek.

REFERENCES

- Baccus, J.R., Baldwin, M.W., & Packer, D.J. (2004). Increasing implicit self-esteem through classical conditioning. *Psychological Science, 15*, 498–502.
- Banaji, M.R., & Bhaskar, R. (2000). Implicit stereotypes and memory: The bounded rationality of social beliefs. In D.L. Schacter & E. Scarry (Eds.), *Memory, brain, and belief* (pp. 139–175). Cambridge, MA: Harvard University Press.
- Campbell, D.T. (1958). Common fate, similarity, and other indices of the status of aggregates of persons as social entities. *Behavioral Science, 3*, 14–25.
- Clark, H.H., & Clark, E.V. (1977). *Psychology and language*. New York: Harcourt Brace Jovanovich.
- Crawford, M.T., Sherman, S.J., & Hamilton, D.L. (2002). Perceived entitativity, stereotype formation, and the interchangeability of group members. *Journal of Personality and Social Psychology, 83*, 1076–1094.
- Gawronski, B., & Bodenhausen, G.V. (2006). Associative and propositional processes in evaluation: An integrative review of implicit and explicit attitude change. *Psychological Bulletin, 132*, 692–731.
- Greenwald, A.G., & Banaji, M.R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review, 102*, 4–27.
- Greenwald, A.G., Banaji, M.R., Rudman, L.A., Farnham, S.D., Nosek, B.A., & Mellott, D.S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review, 109*, 3–25.
- Greenwald, A.G., McGhee, D.E., & Schwartz, J.K.L. (1998). Measuring individual differences in implicit cognition: The Implicit Association Test. *Journal of Personality and Social Psychology, 74*, 1464–1480.
- Greenwald, A.G., 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., Nosek, B.A., Banaji, M.R., & Klauer, K. (2005). Validity of the salience asymmetry interpretation of the Implicit Association Test: Comment on Rothermund and Wentura (2004). *Journal of Experimental Psychology: General, 134*, 420–425.
- Gregg, A.P. (2000). *The hare and the tortoise: The origins and dynamics of implicit and explicit attitudes*. Unpublished doctoral dissertation, Yale University, New Haven, CT.
- Gregg, A.P., Seibt, B., & Banaji, M.A. (2006). Easier done than undone: Asymmetry in the malleability of automatic preferences. *Journal of Personality and Social Psychology, 90*, 1–20.
- Hamilton, D.L., & Gifford, R.K. (1976). Illusory correlation in interpersonal perception: A cognitive basis of stereotypic judgments. *Journal of Experimental Social Psychology, 12*, 392–407.
- Heider, F. (1958). *Interpersonal relations*. New York: Wiley.
- Johnson, M.K., Hashtroudi, S., & Lindsay, D.S. (1993). Source monitoring. *Psychological Bulletin, 114*, 3–28.
- McConnell, A.R., Sherman, S.J., & Hamilton, D.L. (1997). Target entitativity: Implications for information processing about individual and group targets. *Journal of Personality and Social Psychology, 72*, 750–762.
- 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., Greenwald, A.G., & Banaji, M.R. (2005). Understanding and using the Implicit Association Test: II. Method variables and construct validity. *Personality and Social Psychology Bulletin, 31*, 166–180.
- 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.), *Social psychology and the unconscious: The automaticity of higher mental processes* (pp. 265–292). New York: Psychology Press.
- Olson, M.A., & Fazio, R.H. (2001). Implicit attitude formation through classical conditioning. *Psychological Science, 12*, 413–417.
- Reyna, V.F., & Brainerd, C.J. (1995). Fuzzy trace theory: An interim synthesis. *Learning and Individual Differences, 7*, 1–75.
- Rothermund, K., & Wentura, D. (2004). Underlying processes in the Implicit Association Test (IAT): Dissociating salience from associations. *Journal of Experimental Psychology: General, 133*, 139–165.
- Rydell, R.J., McConnell, A.R., Mackie, D.M., & Strain, L.M. (2006). Of two minds: Forming and changing valence-inconsistent implicit and explicit attitudes. *Psychological Science, 17*, 954–958.
- Sloman, S.A. (1996). The empirical case for two systems of reasoning. *Psychological Bulletin, 119*, 3–22.
- Smith, E.R., & DeCoster, J. (2000). Dual-process models in social and cognitive psychology: Conceptual integration and links to underlying memory systems. *Personality and Social Psychology Review, 4*, 108–131.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review, 8*, 220–247.

(RECEIVED 3/20/07; REVISION ACCEPTED 8/7/07)