VU Research Portal

On the inexplicability of the implicit: differences in the information provided by implicit and explicit tests

DeCoster, J.; Banner, M.J.; Smith, E.R.; Semin, G.R.

published in

Social Cognition

2006

DOI (link to publisher)

[10.1521/soco.2006.24.1.5](https://doi.org/10.1521/soco.2006.24.1.5)

document version

Publisher's PDF, also known as Version of record

[Link to publication in VU Research Portal](#)

citation for published version (APA)

DeCoster, J., Banner, M. J., Smith, E. R., & Semin, G. R. (2006). On the inexplicability of the implicit: differences in the information provided by implicit and explicit tests. *Social Cognition*, 24, 5-21.
<https://doi.org/10.1521/soco.2006.24.1.5>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

E-mail address:

vuresearchportal.ub@vu.nl

ON THE INEXPLICABILITY OF THE IMPLICIT: DIFFERENCES IN THE INFORMATION PROVIDED BY IMPLICIT AND EXPLICIT TESTS

Jamie DeCoster
University of Alabama

Michèle J. Banner
Science Applications International Corporation, McLean, VA

Eliot R. Smith
Indiana University

Gün R. Semin
Free University Amsterdam

Implicit measures are often preferred to overt questioning in many areas of psychology. Their covert nature allows them to circumvent conscious expectations and biases, theoretically providing more objective indicators of people's true attitudes and beliefs. However, we argue that implicit and explicit measures tap into different memory systems, so that the interpretation of implicit measures is not as straightforward as the interpretation of explicit measures. We conducted an experiment investigating the relation between implicit and explicit measures of person impressions. The results demonstrate that a single stimulus can have opposite effects on implicit and explicit measures, supporting the theory that the measures reflect the contents of different memory systems. We suggest that implicit measures reflect simple associations stored in a "slow-learning" memory system, while explicit measures reflect a combination of these associations with contextually dependent memories stored in a "fast-binding" memory system.

This article was supported by a grant from the Netherlands Organization for Research (NWO #575-12.020).

Address correspondence to Jamie DeCoster at The University of Alabama, Box 870348, Tuscaloosa, AL 35487-0348; E-mail: jamie@ua.edu.

This article was accepted for publication under the former editor, Donal Carlston.

Whenever you ask someone a question, you can expect that the motivations of your potential informant will have a strong influence on the answer that you get. People who want to please you will change their answer to reflect what they think you want to hear. People who want to impress you with their knowledge will provide you with more information than you need to hear. People who want to deceive you may actually do their best to lead you to the wrong conclusion. While there is always the possibility that the response will actually answer the question you asked, you will typically have to take these and other potential biases into account when interpreting what people say. For this reason, many researchers have come to prefer implicit measures of psychological constructs, where participants' mental states are inferred from their performance on a task, to explicit measures, where participants are directly asked to report their own introspections. Implicit measures are now commonly used in many areas of psychology, including attitudes (Fazio, Sanbonmatsu, & Powell, 1986; Greenwald, McGhee, & Schwartz, 1998), self-esteem (Koole, Dijksterhaus, & van Knippenberg, 2001), prejudice (Dovidio, Kawakami, Johnson, Johnson, & Howard, 1997), and impression formation (Carlston & Skowronski, 1994; Winter & Uleman, 1984).

Recently, however, there has been debate regarding the best way to interpret the responses to implicit measures. Initially implicit measures were developed to be relatively straightforward substitutes for explicit questions whose validity could be compromised by response biases (e.g., Dovidio & Fazio, 1992; Gaertner & McLaughlin, 1983). However, studies indicated that the relations between implicit and explicit measures were often very weak (e.g., Lowery, Hardin, & Sinclair, 2001; Rudman, Ashmore, & Gary, 2001; Wittenbrink, Judd, & Park, 2001). Researchers first assumed that this was because the explicit measures allowed for the possibility of deception, but later investigators considered whether these differences might be caused by something more fundamental. In a review of the use of implicit measures in social psychology, Greenwald and Banaji (1995) suggested a firm distinction between implicit and explicit cognition. They define implicit effects as those that influence participants' thoughts or behaviors in a way that is not detectable by introspective or self-report measures. These ef-

fects appear to operate automatically and outside of conscious awareness. On the other hand, explicit processes require the application of conscious thought, and so could potentially be described by the person doing the processing. Wilson, Lindsey, and Schooler (2000) proposed a “dual attitudes” model, claiming that people have both automatic and conscious components to their attitudes. When people make a deliberate decision as to what they should do, the conscious component of their attitude (corresponding to what is captured by explicit measures) will have a dominating impact on their behaviors. When people act without much deliberation, the automatic component of their attitude (corresponding to what is captured by implicit measures) will have a dominating impact on their behaviors.

It turns out that the field of attitudes is not the only one where there seem to be separate conscious and automatic contributions to behavior. Smith and DeCoster (2000) analyzed such dual-process models across social and cognitive psychology and discovered a large number of similarities across the models. They then presented a general framework to explain how and when people will use automatic and conscious processing. Smith and DeCoster’s formulation is similar to the dual attitudes model in that it proposes that people tend to rely on automatic processing unless they have strong motivations and adequate capacity to process consciously, in which case they will perform conscious processing. However, they go one step further and claim that the differences in the two processing modes arise from the fact that they make use of different types of memories. Smith and DeCoster claim that automatic processing is based purely on a “slow-learning” memory system (believed to physically reside in the neocortex—see McClelland, McNaughton, & O’Reilly, 1995) whose job is to represent long-term, stable knowledge about the environment. Information in this system is not represented by discrete memories or exemplars. Instead, the authors propose that this knowledge is stored as a set of associations among environmental characteristics, similar to the way that a connectionist memory can store information about a large number of stimuli in a single set of weighted links (Smith & DeCoster, 1998).

Conscious processing is also assumed to make use of the slow-learning memory system, but it also adds information from

a separate “fast-binding” memory system (believed to physically reside in the hippocampus—see McClelland, et al., 1995). This latter system is designed to capture the details of specific experiences. Instead of representing information as a set of associations, the fast-binding system is believed to store information about a specific experience as a coherent whole. The representations in this system will contain information about specific events, including elements of the context. The fast-binding system can then use these relations as part of its interpretation of the stimulus.

It is important to note Smith and DeCoster (2000) allow that both the slow-learning and the fast-binding systems can be affected by the exposure to a single stimulus. The difference in their names reflects what the system as a whole is designed represent, rather than the systems’ abilities to respond to new information. The content of the slow-learning system is designed to represent an amalgam of experiences over a long period of time, so the system as a whole will change only slowly over time. The fast-binding system, on the other hand, is designed to store memories of specific episodes, and so the system as a whole needs to be able to change quickly to fully represent new episodes.

The presence of separate but interacting memory systems is consistent with the results of several prior studies. Research by Hastie and Park (1986) demonstrated that people’s memory for the information that led to an evaluative judgment is dissociated from their memory for the evaluation itself. This can be easily explained if the judgment is stored in the slow-learning memory system as an association built up between the target and a positive or negative evaluation, while the data used to form the judgment are kept as verbal statements in the fast-binding memory system. Similarly, research by Martin, Seta, and Crelia (1990) has demonstrated that while diverting conscious attention prevents conscious processing and the formation of verbal memories (those stored in the fast-binding system), it does not influence people’s abilities to form simple associations (those stored in the slow-learning system).

We wish to propose that implicit measures generally reflect the results of automatic processing while explicit measures generally reflect the results of conscious processing. Combined with Smith and DeCoster’s (2000) discussion of the relations between pro-

cessing modes and memory systems, this provides us with a basis for understanding the differences in the information provided by implicit and explicit measures. Specifically, we would claim that implicit measures primarily reflect associations maintained in the slow-learning memory system while explicit measures reflect a combination of the slow-learning memory system with the fast-binding memory system.

One important implication of having two independent memory systems is that the contents of the two systems do not necessarily need to be consistent with each other. Even though the content of both memory systems is ultimately shaped by the same experiences, differences in the way that the two systems process this information could potentially lead to different representations of the same event. The fact that the two systems store their representations in different areas of the brain means that any inconsistencies between them do not have to be resolved. We would next like to present a demonstration that not only can the implicit and explicit measures of parallel constructs be distinct, but also that they can actually contradict each other.

DESIGN AND HYPOTHESES

This study was intended to clarify the differences between how information is represented in the slow-learning and fast-binding memory systems, and in particular demonstrate that a single stimulus event might lead to different (even opposite) representations in the two systems. In their own presentation of a dual-process model, Strack and Deutsch (2004) specifically claim that negations (statements that something is not true, such as "Sam is not messy") are processed differently by automatic and conscious processing. They propose that in conscious processing, the word "not" is used to modify the meaning of the word "messy," while in automatic processing, both "not" and "messy" are associated with Sam individually. We therefore suspected that exposure to a linguistic negation might create a dissociation between the contents of the two memory systems. The fast-binding memory system, reflecting the operation of linguistic and logically based conscious processing, should be able to represent the meaning of

the negation. Thus, a memory representation might be formed that indicates Sam is not messy, or perhaps that Sam is neat. In contrast, the slow-learning memory system deals not with linguistic constructs or with logical operations such as negation, but simply with associations. Encountering a statement that Sam is not messy would lead to the joint activation of the concepts of Sam and messy, creating an association between these two concepts. We therefore suspect that exposure to a negation will lead to the creation of contradictory associations in the two memory systems.

We decided to conduct our investigation in the area of impression formation. Our basic method would be to provide participants with a photograph paired with trait information. This information would state that the person either did or did not have a given trait. We wanted to compare implicit and explicit measures of the association between the person and the stimulus trait, specifically when participants are told that a particular individual does not have a trait. We predict that when asked an explicit question about their impression, participants will indeed report that the target did not have the paired trait, since they will be able to make use of a coherent impression of the target found in their fast-binding system. However, we suspected that an implicit measure would indicate an association between the target and the trait since the slow-learning system would simply record that the two stimuli had been paired together.

METHOD

RESEARCH PARTICIPANTS

Eighty-one participants from introductory psychology classes at Purdue University completed the study in partial fulfillment of course requirements.

PROCEDURE

Each participant was seated in an individual room in front of a computer. They were told that they would be working either on

the computer or in a booklet next to the computer, and that all instructions would appear on the computer screen.

Participants were first asked to examine a number of person descriptions, each consisting of a photograph and a sentence indicating that the individual either did or did not possess a particular trait (e.g., *Rolanda IS helpful*; *Sam is NOT neat*). When each pairing was presented, participants were instructed to copy down the sentence from the screen into the provided booklet. On the next line, they were asked to come up with their own example of a behavior that the pictured person might perform. Participants completed a total of 16 behavior-writing trials, each including either a positive or negative trait, preceded or not preceded by the word “not.” Participants never saw both traits of a diametrically opposed pair (such as “neat” and “messy”). Participants next performed a distractor task for approximately five minutes (naming countries beginning with each letter of the alphabet), supposedly as a measure of long-term memory. Participants then completed either an implicit or explicit measure testing the associations created in the initial exposure task between the traits and photographs.

The procedure for the remainder of the study varied depending on whether the participant was in the explicit measure condition or the implicit measure condition. For the explicit measure, participants were directly told that we wanted to collect their impressions of the individuals they saw in the exposure task. In each of 16 trials, participants viewed a photograph that they had seen in the exposure task along with a set of three traits. They were asked to rate the extent to which they felt the person in the photo possessed each of the traits on a scale of one to seven, where higher numbers indicated that the trait was more applicable. Two of the traits in each trial were unrelated to those presented in the person descriptions. Half of the time the remaining trait was the one originally paired with the photo in the person descriptions (but without the word “not” in all cases), while in the other half the trait was the bipolar opposite of the original trait. Whether the critical trait was rated first, second, or third was randomly determined on each trial.

An illustration of how the explicit measure was implemented is provided in Figure 1.

FIGURE 1. Conditions in the explicit ratings measure.

Each participant performing the explicit rating task had (A) four trials where they wrote sentences saying that a person had trait X in the exposure task and then evaluated the person on trait X in the rating task, (B) four trials where they wrote sentences saying that a person had trait X in the exposure task and then evaluated the person on the bipolar opposite of trait X in the rating task, (C) four trials where they wrote sentences saying that a person did NOT have trait X in the exposure task and then evaluated the person on trait X in the rating task, and (D) four trials where they wrote sentences saying that a person did NOT have trait X in the exposure task and then evaluated the person on the bipolar opposite of trait X in the rating task. In this task, higher ratings indicate a stronger explicit association. We therefore expect the ratings provided in conditions A and D (where the trait being rated is logically consistent with the information provided in the initial exposure) to be higher than the ratings provided in conditions B and C.

We used Carlston and Skowronski's (1994) "savings in relearning" procedure as our implicit measure. This method is based on the work of Ebbinghaus (1885/1964), who demonstrated that it was easier to relearn a pairing between two objects when the two objects had been paired together previously. Recall that in the exposure task, participants were asked to write a total of 16 sentences, 8 pairing photos with statements that the pictured individual possessed a trait, and 8 pairing photos with statements that the pictured individual did not possess a trait. Participants completing the implicit measure were then asked to memorize 30 photo-trait pairs. The instructions indicated that some of the photographs might have been viewed earlier in the study. Eight paired a photo and trait that were originally presented together in the person descriptions (but without the word "NOT" in all cases). Eight paired a photo presented in the person descriptions with the bipolar opposite of the trait presented with that photo. In each condition, half of the traits were evaluatively positive and half were evaluatively negative. The remaining 14 pairs included novel photographs and traits. Participants then completed a second distractor task where they were asked to perform complex mental arithmetic problems for approximately five minutes. Finally, they performed a cued-recall task where a photo was displayed on the screen for eight seconds and participants were asked to write down the trait that had been paired with the photo in the memorization task. Each photo only appeared once in the memory task, and only one member of each bipolar trait pair appeared in the memory task.

An illustration of how the implicit measure was implemented is provided in Figure 2. Each participant performing the implicit savings in relearning task had (A) four trials where they wrote sentences saying that a person had trait X in the exposure task and were then asked to memorize a pairing of that person with trait X in the memory task, (B) four trials where they wrote sentences saying that a person had trait X in the exposure task and were then asked to memorize a pairing of that person with the bipolar opposite of trait X in the memory task, (C) four trials where they wrote sentences saying that a person did NOT have trait X in the exposure task and were then asked to memorize a pairing of that person with trait X in the memory task, and (D) four trials where

FIGURE 2. Conditions in the implicit savings in relearning measure.

they wrote sentences saying that a person did NOT have trait X in the exposure task and were then asked to memorize a pairing of that person with the bipolar opposite of that trait in the memory task. The memory task also included 14 additional pairings of novel photos with traits unrelated to those seen in the exposure task. In this task, better memory indicates a stronger implicit association. We therefore expect the proportion of photo–trait pairs correctly recalled in conditions A and B (where same trait is used in both the exposure task and the memory task) to be higher than the ratings provided in conditions C and D.

We have several reasons to believe that the savings in relearning method provides an implicit measure of the association between two objects. Evidence for this is provided by the fact that the savings effect is present long after explicit memory for the original pairing is long forgotten (Ebbinghaus, 1885/1964). Carlston and Skowronski (1994) have also demonstrated that this measure could be used to detect trait inferences automatically and unintentionally formed when reading behavioral statements.

TABLE 1. Explicit Trait Ratings by Exposure Condition and Response Match

	Exposure "IS trait"	Exposure "Is NOT trait"
Rating trait same as exposure trait	A = 4.46 (1.22)	B = 3.67 (1.06)
Rating trait opposite of exposure trait	C = 3.61 (0.85)	D = 4.51 (0.89)

Note. The letters A through D correspond to the conditions as presented in Figure 1. The number following the equal sign is the mean rating for the condition, while the number in parentheses is the standard deviation.

Although our procedure differs, the rationale for our predictions is identical.

RESULTS

There were no main effects or interactions with trait valence on either of the two tasks, so all of the results reported below average over both positive and negative traits.

EXPLICIT RATING TASK

The explicit ratings were submitted to a within-subjects ANOVA including exposure condition (whether the target was said to possess the trait or not at exposure) and trait match (whether the ratings were made of the same trait used at exposure or its bipolar opposite) as factors. Table 1 presents the explicit ratings for each cell in the design. There were no main effects of exposure condition $F(1, 39) = 0.01, p > .9$, or trait match $F(1, 39) = 0.16, p > .6$, on the rating. However, we did observe a significant interaction between exposure condition and trait match $F(1, 39) = 20.53, p < .0001$. As predicted, the ratings found in cells A and D of Table 1 are significantly higher than those in cells B and C, $t(39) = 3.84, p < .0005$. When told that a target possessed a trait, explicit ratings were higher for that trait than its bipolar opposite. When told that a target did not possess a trait, the explicit rating of the bipolar opposite was stronger. This indicates that participants accurately used the "not" prefix when forming their explicit judgments of the targets.

TABLE 2. Proportion Correct by Exposure Condition and Response Match

	Exposur "IS trait"	Exposure "Is NOT trait"
Recall trait same as exposure trait	A = .470 (.336)	B = .427 (.302)
Recall trait opposite of exposure trait	C = .323 (.264)	D = .329 (.308)

Note. The letters A through D correspond to the conditions as presented in Figure 2. The number following the equal sign is the mean proportion correct for the condition, while the number in parentheses is the standard deviation.

IMPLICIT SAVINGS IN RELEARNING TASK

We then submitted the recall accuracies of the implicit task to this same analysis. Table 2 presents the recall accuracies for each cell in the design. We observed a significant main effect of trait match $F(1, 40) = 9.82, p < .004$, but no influence of the exposure condition $F(1, 40) = 0.21, p > .6$. The interaction of trait match and exposure condition was also non-significant $F(1, 40) = 0.40, p > .5$. As predicted, the proportion correct in cells A and B of Table 2 is significantly higher than the proportion correct in cells C and D $t(40) = 3.134, p < .004$. Implicit associations were strongest between the target and the exact trait presented at exposure, regardless of whether the target was said to possess the trait or not. This indicates that the "not" prefix did not have a significant influence on the associations detected by the implicit measure.

While it is possible that savings could result from the use of some explicit strategy, we do not believe that that is the case in this study. First, participants did not expect that their impressions of the individuals in the first part of the study would be examined later on, and we included a difficult distractor task between the exposure and test phases to make using such strategies more difficult. Second, participants might have recalled inconsistent photo-trait pairings because this information is sometimes well-remembered. But in that case, traits in condition D should have been recalled as well as traits in condition B. Finally, even if some participants did use an explicit strategy to assist their memory, it would simply make it more difficult for us to detect differences in the impressions tapped by implicit and explicit

measures, so this cannot act as an alternative explanation for any differences we do observe.

DISCUSSION

The results from both the explicit and implicit measures support our predictions, and are consistent with the interpretation that they reflect the contents of different memory systems. Our study examined how providing people with information that a target individual did not possess a trait would affect their impressions reported by implicit and explicit measures. We found that people asked to explicitly report their impression stated that the target lacked the trait. However, people whose impressions were measured implicitly showed an association between the target and the denied trait. This illustrates that not only can there be differences between the results of implicit and explicit measures of person impressions, but that they can actually contradict each other.

Although differences between implicit and explicit measures have commonly been ascribed to intentional correction effects (Dovidio & Fazio, 1992), we do not believe that this is the case with our study. This effect was found across a collection of 16 different traits covering a broad set of characteristics, so it is unlikely to be caused by attempts to correct for stereotyping. Additionally, the same effect was found for positive and negative traits, so it is unlikely that our results are a result of people trying to present more positive or more negative impressions. Since there appears to be no reason to believe that our results are due to a correction effect, we believe that they indicate that people have formed two separate impressions of the target person, and that they coexist in memory even though they contradict each other. A parsimonious explanation for this is provided by the dual-memory system discussed above. It would claim that the explicit measure reflects a representation in the fast-binding system while the implicit measure reflects a representation in the slow-learning memory system. The fact that the representations for these systems are stored in different areas of the brain allows them to independently coexist. While it is certainly possible for implicit and explicit measures to provide similar results, they can differ significantly if the ex-

implicit measure makes strong use of information contained in the fast-binding memory system and this information differs from that contained in the slow-learning memory system.

IMPLICATIONS OF A DUAL-MEMORY SYSTEM INTERPRETATION

From a dual-memory system perspective, we would expect that measures that are processed automatically solely reflect the contents of the slow-learning system, while measures that are processed consciously can reflect either the contents of the slow-learning or the fast-binding system. Our results that the memories tapped by implicit and explicit measures of person impressions appear to reflect distinct mental representations provides evidence that distinct memory systems are implicated in the processing of these two measures. We would also like to suggest that the dual-memory system interpretation also applies to implicit and explicit measures of other phenomena, assuming that the implicit measure represents automatic processing while the explicit measure represents conscious processing.

So what does this perspective have to say about associations detected using implicit measures? Does it mean that a person who demonstrates an association between two objects on an implicit measure consciously believes that the two objects are related? Clearly not, since implicit measures often have only a small relation to their explicit counterparts (Dovidio, Kawakami, & Gaertner, 2002). Does it mean that the person has some "gut feeling" that the two objects are related that can't easily be put into words? Not necessarily. In our study, the explicit measure of the impression simply asks participants to rate the extent to which they thought that a target person had particular traits, without asking them for any explanation or justification for their rating. We would therefore expect that our explicit ratings would also have reflected any vague feelings that the target person and trait were related, even if they could not be verbalized. Given that we found divergent effects on our implicit and explicit measures, implicit measures must at least in part reflect aspects of memory that have no direct influence on explicit ratings.

The dual-memory system account for implicit measures is more complicated than either of these interpretations. It would claim that implicit measures reflect associations within the slow-learning memory system. Observing an association between two things on an implicit measure would therefore indicate that those two things had been paired together in the person's past experience. This could represent either a recent pairing (as in our study) or a pairing that had been consistently observed over time. It would not necessarily reflect any beliefs, conscious or unconscious, held by the person in which the association is measured.

We must acknowledge two important limitations of our study. First, the implicit and explicit measures that we used did differ substantially. We cannot make any claims regarding the relative strength of the implicit and explicit associations formed in our study since our measures are clearly not equally sensitive. It could be useful to replicate this study using implicit and explicit measures that parallel each other more closely. In addition, we examined different participants in the implicit and explicit conditions. Our claims that implicit and explicit measures tap different memory systems would be even more compelling if we could demonstrate dissociations between the two types of measures on the same participants.

Our research supports the idea that implicit measures tap into the simple associations stored in the slow-learning memory system, so increasing our understanding of these associations will better enable us to interpret these measures. We demonstrate two important properties of this associative system. First, the contents of this system simply represent what elements have been paired together in the environment and may therefore fail to capture inferences and conclusions deriving from conscious processing of the events. Second, information in the slow-learning memory system appears to be stored separately from information in the fast-binding memory system, so we should expect that the two might be inconsistent. Our success in explaining the results in this study using a dual-memory systems model suggests that further research on this model will help us understand how the characteristics of our measures affect the results they provide.

REFERENCES

- Carlston, D. E., & Skowronski, J. J. (1994). Savings in the relearning of trait information as evidence for spontaneous inference generation. *Journal of Personality and Social Psychology*, 66, 840–856.
- Dovidio, J. F., & Fazio, R. H. (1992). New technologies for the direct and indirect assessment of attitudes. In J. M. Tanur (Ed.), *Questions about questions: Inquiries into the cognitive bases of surveys*. New York, NY: Sage.
- Dovidio, J. F., Kawakami, K., Johnson, C., Johnson, B., & Howard, A. (1997). On the nature of prejudice: automatic and controlled processes. *Journal of Experimental Social Psychology*, 33, 510–540.
- Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial relations. *Journal of Personality and Social Psychology*, 82, 62–68.
- Ebbinghaus, H. (1964). *Memory: A contribution to experimental psychology*. New York: Dover. (Original work published 1885)
- Fazio, R. H., Sanbonmatsu, D. M., & Powell, M. C. (1986). On the automatic activation of attitudes. *Journal of Personality and Social Psychology*, 50, 229–238.
- Gaertner, S. L., & McLaughlin, J. P. (1983). Racial Stereotypes: associations and ascriptions of positive and negative characteristics. *Social Psychology Quarterly*, 46, 23–30.
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: attitudes, self-esteem, and stereotypes. *Psychological Review*, 102, 4–27.
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: the implicit association test. *Journal of Personality and Social Psychology*, 74, 1464–1480.
- Hastie, R., & Park, B. (1986). The relationship between memory and judgment depends on whether the judgment task is memory-based or on-line. *Psychological Review*, 93, 258–268.
- Koole, S. L., Dijksterhaus, A., van Knippenberg, A. (2001). What's in a name: implicit self-esteem and the automatic self. *Journal of Personality and Social Psychology*, 80, 669–685.
- Lowery, B. S., Hardin, C. D., & Sinclair, S. (2001). Social influence effects on automatic racial prejudice. *Journal of Personality and Social Psychology*, 81, 842–855.
- Martin, L. L., Seta, J. J., & Crelia, R. A. (1990). Assimilation and contrast as a function of people's willingness and ability to expend effort in forming an impression. *Journal of Personality and Social Psychology*, 59, 27–37.
- McClelland, J. L., McNaughton, B. L., & O'Reilly, R. C. (1995). Why there are complementary learning systems in the hippocampus and neocortex: Insights from the successes and failures of connectionist models of learning and memory. *Psychological Review*, 102, 419–457.
- Rudman, L. A., Ashmore, R. D., & Gary, M. L. (2001). "Unlearning" automatic biases: The malleability of implicit prejudice and stereotypes. *Journal of Personality and Social Psychology*, 81, 856–868.
- Smith, E. R., & DeCoster, J. (1998). Knowledge acquisition, accessibility, and use

- in person perception and stereotyping: Simulation with a recurrent connectionist network. *Journal of Personality and Social Psychology*, 74, 21–35.
- 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.
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101–126.
- Winter, L., & Uleman, J. S. (1984). When are social judgments made? Evidence for the spontaneousness of trait inferences. *Journal of Personality and Social Psychology*, 47, 237–232 (See correction in *Journal of Personality and Social Psychology*, 1986, 50, 355).
- Wittenbrink, 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.