

How Should We Think About Implicit Measures and Their Empirical “Anomalies”?

Bertram Gawronski

Department of Psychology, University of Texas at Austin

Michael Brownstein

Department of Philosophy, John Jay College of Criminal Justice/City University of New York

Alex Madva

Department of Philosophy, California State Polytechnic University Pomona

Based on a review of several “anomalies” in research using implicit measures, Machery (2021) dismisses the modal interpretation of participant responses on implicit measures and, by extension, the value of implicit measures. We argue that the reviewed findings are anomalies only for specific—influential but long-contested—accounts that treat responses on implicit measures as uncontaminated indicators of trait-like unconscious representations that coexist with functionally independent conscious representations. However, the reviewed findings are to-be-expected “normalities” when viewed from the perspective of long-standing alternative frameworks that treat responses on implicit measures as the product of dynamic processes that operate on momentarily activated, consciously accessible information. Thus, although we agree with Machery that the modal view is empirically unsupported, we argue that implicit measures can make a valuable contribution to understanding the complexities of human behavior if they are used wisely in a way that acknowledges what they can and cannot do.

Keywords: anomalies; consciousness; implicit measures; mental representation; reliability

In response to our descriptive review of ongoing debates about what implicit measures measure (Brownstein, Madva, & Gawronski, 2019), Machery (2021) discusses several important findings in research using implicit measures that he describes as “anomalies.” Based on the identified anomalies, Machery dismisses both the modal paradigm in research using implicit measures and, by extension, the value of implicit measures. Here, we respond to Machery’s critique, arguing that the findings reviewed by Machery are anomalies only for specific—influential but long-contested—accounts that treat responses on implicit measures as uncontaminated indicators of trait-like unconscious representations that coexist with functionally independent conscious representations. However, the reviewed findings are to-be-expected “normalities” when viewed from the perspective of long-standing alternative frameworks that treat responses on implicit measures as the product of dynamic processes that operate on momentarily activated, consciously accessible information. Thus, although we agree with Machery that the modal view is theoretically, empirically, and methodologically unsubstantiated, its inconsistency with the available evidence does not imply that implicit measures are useless for understanding the underpinnings of human behavior.

We focus on this central disagreement in our reply. Before explaining our position, however, we briefly note a few important issues which we will set aside for now. We agree with Machery that models of the cumulative effects of social biases require empirical support (Mallon, 2021) and that the effectiveness of

implicit-bias training requires scientific scrutiny (Brownstein, Madva, & Gawronski, 2020; Madva, 2020). We also agree that poor science communication is a fundamental problem in this area, although it is crucial to clearly distinguish between problems in scientific research and problems in the public communication of scientific research. While extremely important, we also set aside debates about the use of implicit measures to investigate macro-level phenomena (e.g., the bias of states and nations; see Payne, Vuletic, & Lundberg, 2017). Finally, we note that Machery’s repeated invocation of how much remains unsettled “30 years after the implicit revolution” (p. 5) takes it as self-evident that 30 years is too long for a scientific research program to settle on the kinds of debates and frameworks we discuss here. This is not at all self-evident. Basic theorizing about emotion, perception, attention, and consciousness is multipronged, conflicted, and ongoing—appropriately so—after much more than 30 years.

In explaining our view on the anomalies identified by Machery, we adopt the same meta-theoretical stance that we used in our review, treating responses on implicit measures as behaviors rather than direct indicators of mental constructs (see De Houwer, Gawronski, & Barnes-Holmes, 2013). More specifically, we understand them as responses to stimuli under suboptimal processing conditions, elicited without intention and with little cognitive elaboration (see De Houwer, Teige-Mocigemba, Spruyt, & Moors, 2009). As such, they can be distinguished from responses on explicit measures, which tend to be shaped by intentions and greater cognitive elaboration. This

theoretically agnostic conceptualization raises two important questions: (1) what underlies responses on implicit measures and (2) in which sense are their underpinnings similar or different from those of responses on explicit measures?

Underlying Representations

From the perspective of our behavioral conceptualization, two important arguments in Machery's critique are that (1) the mental representations underlying responses on implicit measures are not unconscious and (2) there is no evidence for the idea that responses on implicit and explicit measures are driven by two distinct underlying representations (e.g., implicit vs. explicit attitudes). We fully agree with these conclusions, which are perfectly consistent with well-established theories in this area (e.g., Fazio, 2007, Gawronski & Bodenhausen, 2006; Strack & Deutsch, 2004). These theories assume that the same associative representations underlie responses on both implicit and explicit measures and that these representations are consciously accessible. Whether the assumption of associative representation is defensible in light of recent evidence is a different question, because a growing body of evidence suggests that the representations underlying responses on implicit and explicit measures have propositional structure (for reviews, see Corneille & Stahl, 2019; De Houwer, Van Dessel, & Moran, 2020; Kurdi & Dunham, 2020). Although research on this question is ongoing, the critical point is that evidence against unconsciousness and dual representations are anomalies only for frameworks that postulate two distinct underlying representations, one being conscious and one being unconscious. However, these assumptions are rejected by alternative frameworks that have shaped the field for a long time.

Reliability

Although we share Machery's general concern about the reliability of implicit measures, his analysis misses important nuances, especially the difference between *internal consistency* and *test-retest stability*. While the internal consistency of a measure sets an upper limit for the proportion of systematic construct variance, the test-retest stability of a measure specifies the proportion of systematic variance that is stable over time. Machery questions the construct validity of implicit measures based on evidence for low test-retest stability, which conflates internal consistency and test-retest stability. A more nuanced analysis suggests that the internal consistency of some implicit measures (e.g., implicit association test; affect misattribution procedure) is quite high, meeting the psychometric standards applied to explicit measures (Gawronski & De Houwer, 2014). For other implicit measures,

internal consistency is only moderate (e.g., go/no-go association task) or unsatisfactory (e.g., evaluative priming task). These differences are important, because the internal consistency of a given measure sets an upper limit for its correlations with other measures, including correlations with itself at a different time. Mathematically, a measure with low internal consistency cannot show high temporal stability, because low internal consistency implies a low signal-to-noise ratio in measurement scores (unless random error is controlled in latent-variable analyses). Yet, high internal consistency does not imply high temporal stability, which is reflected in the finding that implicit measures with high internal consistency nevertheless show only moderate-to-low temporal stability (Gawronski, Morrison, Phillips, & Galdi, 2017).

Does this mean that these measures have low construct validity? If the measure is supposed to capture a temporally stable, trait-like construct, the answer would be *yes*. However, such a conclusion would be ill-founded for measures of state-like constructs that are known to be highly context-sensitive. Indeed, responses on implicit measures have long been known to be highly context-sensitive (Blair, 2002; Gawronski & Sritharan, 2010) and research using latent-state-trait modeling suggests that the total amount of systematic construct variance on implicit measures comprises a mix of situation-related and person-related variance (e.g., Dentale, Veccione, Ghezzi, & Barbaranelli, 2019; Koch, Ortner, Eid, Caspers, & Schmitt, 2014; Lemmer, Gollwitzer, & Banse, 2015; Schmukle & Egloff, 2005). Although these findings conflict with the idea that responses on implicit measures exclusively reflect trait-like constructs that are highly stable over time, they are consistent with theories that treat responses on implicit measures as the product of dynamic processes that operate on momentarily accessible information (e.g., Cunningham, Zelazo, Packer, & Van Bavel, 2007; Ferguson & Bargh, 2003; Gawronski & Bodenhausen, 2006; Schwarz, 2007). Because incidental changes in contexts over time are unlikely to be identical for different participants, temporal stability should be low according to these theories. This does not mean that the measures have low construct validity, but that the measured states have changed in different ways for different participants in the sample. By conflating internal consistency and test-retest stability, Machery's analysis misses this important nuance.

Behavior Prediction

Because the states underlying responses on implicit measures fluctuate over time, implicit measures should perform poorly in predicting other behavioral criteria when there is a delay between the completion of the implicit measure and the measurement of the to-be-predicted behavior (see Gawronski, 2019). From this

perspective, the relatively weak relations between implicit measures and behavioral criteria obtained in meta-analyses (Cameron, Brown-Iannuzzi, & Payne, 2012; Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Kurdi et al., 2019; Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013) are not surprising, especially for cases where there is a delay between the implicit measure and the measurement of the behavioral criterion. This is not a post-hoc rationalization of inconvenient findings, but a straightforward prediction of theories that treat responses on implicit measures as the product of dynamic processes that operate on momentarily accessible information. Poor predictive validity over time is certainly disappointing for everyone who was hoping to improve the prediction of future behavior by using implicit measures, but it constitutes an anomaly only for theories claiming that implicit measures capture temporally stable traits.

A different question concerns the prediction of behavior without delays within the same context. One does not have to be a proponent of dual-process theories to agree with the methodological dictum that relations between a predictor measure and a behavioral-criterion measure should increase as a function of increasing convergence of the mechanisms underlying responses on the two measures and the processing conditions during measurement (see Gawronski & De Houwer, 2014). Thus, even from a theoretically agnostic view, implicit measures should be superior in predicting unintentional behavior resulting from low elaboration, while explicit measures should be superior in predicting intentional behavior resulting from high elaboration. Machery notes that this hypothesis received little support in meta-analyses on the predictive validity of implicit measures.

It is correct that little evidence supports the predicted dissociation pattern in analyses based on post-hoc categorizations of behavioral criteria in terms of processing features, comparing predictive relations between studies. However, there is substantial evidence for the predicted dissociation pattern in analyses based on original categorizations of behavioral criteria in a given study, comparing predictive relations within studies (Cameron et al., 2012). Machery dismisses the latter findings, because the studies were not preregistered. However, the same is true for the overwhelming majority of studies cited by Machery in support of his arguments. It is unconvincing to selectively invoke concerns about lack of preregistration only for findings inconsistent with a preferred narrative.

Furthermore, while explicit measures have shown relatively strong relations with deliberate behavior and relatively weak relations with spontaneous behavior in meta-analytic comparisons between studies, implicit measures have shown small positive relations with both

spontaneous and deliberate behavior. To the extent the dissociation hypothesis is correct, this asymmetric pattern can result from simple differences in the reliability of behavioral-criterion measures, given that measures of spontaneous behavior are often based on a single observation whereas measures of deliberate behavior tend to aggregate multiple observations, thereby controlling for measurement error. As explained by Gawronski (2019, p. 582):

For explicit measures, the described asymmetry in the reliability of behavioral criteria should produce strong relations to deliberate behavior (because of matching processing conditions with a reliable behavioral criterion) and relatively weak or nonsignificant relations to spontaneous behavior (because of mismatching processing conditions with an unreliable behavioral criterion). In contrast, for implicit measures, the described asymmetry in the reliability of the behavioral criteria should produce relatively weak relations to both spontaneous behavior (because of low reliability of the behavioral measure) and deliberate behavior (because of mismatching processing conditions).

Because differences in the reliability of behavioral criteria are more likely to be controlled in studies designed to test the hypothesized dissociation, these considerations explain why the full dissociation pattern emerged only in meta-analytic comparisons of behavioral criteria within studies, but not in meta-analytic comparisons of behavioral criteria between studies (Cameron et al., 2012). This conclusion is also consistent with a large body of research showing the hypothesized dissociation pattern for predictions of the same behavior under different processing conditions (e.g., prediction of a given behavior under conditions of high vs. low cognitive load; prediction of a given behavior for people with intuitive vs. deliberative thinking styles; for a review, see Friese, Hofmann, & Schmitt, 2008). The latter findings cannot be dismissed as the product of arbitrary post-hoc classifications of the behavioral criteria, because the to-be-predicted behavioral criterion is the same and there is no ambiguity about the classifications of the theoretically relevant moderators (e.g., high vs. low cognitive load; intuitive vs. deliberative thinking style). Thus, although meta-analytic findings regarding the predicted dissociation pattern are mixed, a more nuanced analysis provides a simple explanation for why the predicted pattern is more likely to emerge in meta-analytic comparisons within studies than comparisons between studies.

Causality

A final “anomaly” identified by Machery is that experimentally induced changes on implicit measures

do not seem to mediate corresponding changes in behavior (Forscher et al., 2019). According to Machery, this meta-analytic finding poses a major challenge to the idea that whatever is being measured by implicit measures is causally effective in influencing downstream behavior. We agree that this finding might seem puzzling, but the expectation of mediation in such cases is based on two questionable premises.

First, it is worth noting that there are no process-pure psychological measures that exclusively reflect a single construct. In research using implicit measures, acknowledgement of this issue has inspired the use of formal modeling approaches to quantify the contributions of multiple distinct processes to responses on implicit measures (for reviews, see Calanchini, 2020; Sherman, Klauer, & Allen, 2010). Moreover, to the extent that all implicit measures have more than one source of systematic variance, observed changes on implicit measures should generalize to behavioral-criterion measures only if the affected source of variance is also relevant for the behavioral criterion (e.g., when high cognitive load affects both performance on the implicit measure and the behavioral criterion). To the extent that responses on the implicit measure are affected via a source of variance that is irrelevant for the behavioral criterion, changes on the implicit measure should not be associated with corresponding changes on the behavioral-criterion measure. Such findings are surprising only when implicit measures are treated as golden pipelines into a single underlying construct, but not when their lack of process-purity is taken into account.

Second, the expectation of a mediating link between changes on implicit measures and changes in behavioral-criterion measures is based on a flawed conceptualization of mediation that has been criticized for a long time (e.g., Bullock, Green, & Ha, 2010; De Houwer et al., 2013; Fiedler, Schott, & Meiser, 2011; Spencer, Zanna, & Fong, 2005). For the current analysis, the most significant issue is that responses on implicit measures are behaviors, and these behaviors should not be equated with their underlying mental constructs (see De Houwer et al., 2013; Eagly & Chaiken, 2007; Gawronski & Bodenhausen, 2015). In this sense, their ontological status is not different from the one of to-be-predicted behaviors. They are just two kinds of behaviors. Although it makes sense to assume that change in a shared underlying mental construct should lead to corresponding changes in both, it makes less sense to assume that changes in one type of behavior is, just as such, causally effective in producing change in a different type of behavior. If the two are treated as two kinds of behaviors, the situation is actually more similar to a common-third-variable scenario, where correlations between two behavioral measures are driven by their shared relation to a

common causal factor that influences both measures. In this case, changes in the underlying causal force can be expected to produce changes in two behavioral measures that are influenced by that factor. However, it would be absurd to treat one of the two behavioral variables as a mediator of changes in the other. These limitations and various other problems of mediation analyses have been discussed extensively in the past 15 years, and they have led some journal editors in psychology to explicitly discourage the use of mediation analyses. Yet, despite the insights gained in these debates, the finding that changes in implicit measures do not mediate changes in behavioral-criterion measures has been mistakenly interpreted as an “anomaly” that would question the validity of implicit measures. As with the other findings reviewed by Machery, this so-called “anomaly” turns into a to-be-expected “normality” when the current state of the field is taken into account.

Individual Differences

An important question in response to our analysis is whether a dynamic-process view on implicit measures involves a notion of radical situationism that puts it at odds with the strong emphasis on individual differences in the modal view. Although dynamic theories tend to focus more heavily on the impact of situation-related compared to person-related factors, they are consistent with influences of both factors, because momentary activation of information in memory is a joint product of both (Cunningham et al., 2007; Gawronski & Bodenhausen, 2006). The same stimulus can activate different thoughts in a given person when the stimulus is encountered in different contexts, suggesting a major role of situation-related factors. Meanwhile, when a given stimulus is encountered within the same context, the stimulus can activate different thoughts in different people, suggesting a major role of person-related factors. In fact, the available evidence suggests that responses on implicit measures are best understood as an interactive product of the two, in that (1) people show systematic differences in their responses to a particular stimulus *within* a given situation, but (2) individual differences within a given situation may not generalize *across* situations (Gawronski & Bodenhausen, 2017). Consistent with this idea, the relatively low test-retest stability of implicit measures has been found to increase to a level that is consistent with the notion of temporally stable traits when the measure included context-stimuli that were meaningfully related to the target stimuli (Gschwendner, Hofmann, & Schmitt, 2008; for related findings, see Cooley & Payne, 2017). Counter to Machery’s dismissal of implicit measures as tools to capture individual differences, these findings suggest that implicit measures capture meaningful differences

between individuals within a given context, and that individual differences measured within a given context are stable over time. If responses on implicit measure were an exclusive product of person-related factors without any impact of situation-related factors (as suggested by the modal view), their temporal ability should be high regardless of the context. Conversely, if responses on implicit measures were the exclusive product of situation-related factors without any impact of person-related factors (as suggested by a radical situationist view), differences between individuals within the same context should be nothing more than measurement noise, which should reduce test-retest correlations when the context is held constant across measurement times. Counter to the two views, implicit measures show greater temporal stability when the context is the same for everyone in the sample and the context is held constant over time, highlighting the significance of person-by-situation interactions.

Conclusion

The claims made here are not post-hoc attempts to explain away “anomalies,” but essential components of well-established frameworks that have shaped the science of implicit measures for nearly two decades. Dismissing implicit measures because they do not provide process-pure indicators of temporally stable, trait-like unconscious representations is like throwing out all the screw drivers in one’s tool box because they are not suitable for getting a nail into the wall. Understanding the complexities of human behavior requires multiple tools, and implicit measures can make a valuable contribution to this endeavor if they are used wisely in a way that acknowledges what they can and cannot do.

References

- Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, *6*, 242-261.
- Brownstein, M., Madva, A., & Gawronski, B. (2019). What do implicit measures measure? *Wiley Interdisciplinary Reviews: Cognitive Science*, *10*(5):e1501.
- Brownstein, M., Madva, A., & Gawronski, B. (2020). Understanding implicit bias: Putting the criticism into perspective. *Pacific Philosophical Quarterly*, *101*, 276-307.
- Bullock, J. G., Green, D. P., & Ha, S. E. (2010). Yes, but what’s the mechanism? (don’t expect an easy answer). *Journal of Personality and Social Psychology*, *98*, 550-558.
- Calanchini, J. (2020). How multinomial processing trees have advanced, and can continue to advance, research using implicit measures. *Social Cognition*, *38*, s165-s186.
- Cameron, C. D., Brown-Iannuzzi, J., & Payne, B. K. (2012). Sequential priming measures of implicit social cognition: A meta-analysis of associations with behaviors and explicit attitudes. *Personality and Social Psychology Review*, *16*, 330-350.
- Cooley, E., & Payne, B. K. (2017). Using groups to measure intergroup prejudice. *Personality and Social Psychology Bulletin*, *43*, 46-59.
- Corneille, O., & Stahl, C. (2019). Associative attitude learning: A closer look at evidence and how it relates to attitude model. *Personality and Social Psychology Review*, *23*, 161-189.
- Cunningham, W. A., Zelazo, P. D., Packer, D. J., & Van Bavel, J. J. (2007) The iterative reprocessing model: A multilevel framework for attitudes and evaluation. *Social Cognition*, *25*, 736-760.
- De Houwer, J., Gawronski, B., & Barnes-Holmes, D. (2013). A functional-cognitive framework for attitude research. *European Review of Social Psychology*, *24*, 252-287.
- De Houwer, J., Teige-Mocigemba, S., Spruyt, A., & Moors, A. (2009). Implicit measures: A normative analysis and review. *Psychological Bulletin*, *135*, 347-368.
- De Houwer, J., Van Dessel, P., & Moran, T. (2020). Attitudes beyond associations: On the role of propositional representations in stimulus evaluation. *Advances in Experimental Social Psychology*, *61*, 127-183.
- Dentale, F., Vecchione, M., Ghezzi, V., & Barbaranelli, C. (2019). Applying the latent state-trait analysis to decompose state, trait, and error components of the self-esteem Implicit Association Test. *European Journal of Psychological Assessment*, *35*, 78-85.
- Eagly, A. H., & Chaiken, S. (2007). The advantages of an inclusive definition of attitude. *Social Cognition*, *25*, 582-602.
- Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition*, *25*, 603-637.
- Ferguson, M. J., & Bargh, J. A. (2003). *The constructive nature of automatic evaluation*. In J. Musch & K. C. Klauer (Eds.), *The psychology of evaluation: Affective processes in cognition and emotion* (pp. 169-188). Mahwah, NJ: Erlbaum.
- Fiedler, K., Schott, M., & Meiser, T. (2011). What mediation analysis can (not) do. *Journal of Experimental Social Psychology*, *47*, 1231-1236.
- Friese, M., Hofmann, W., & Schmitt, M. (2008). When and why do implicit measures predict behavior? Empirical evidence for the moderating role of opportunity, motivation, and process reliance. *European Review of Social Psychology*, *19*, 285-338.
- Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019).

- A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology*, *117*, 522–559.
- Gawronski, B. (2019). Six lessons for a cogent science of implicit bias and its criticism. *Perspectives on Psychological Science*, *14*, 574–595.
- 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.
- Gawronski, B., & Bodenhausen, G. V. (2015). Social-cognitive theories. In B. Gawronski, & G. V. Bodenhausen (Eds.), *Theory and explanation in social psychology* (pp. 65–83). New York: Guilford Press.
- Gawronski, B., & Bodenhausen, G. V. (2017). Beyond persons and situations: An interactionist approach to understanding implicit bias. *Psychological Inquiry*, *28*, 268–272.
- Gawronski, B., & De Houwer, J. (2014). Implicit measures in social and personality psychology. In H. T. Reis, & C. M. Judd (Eds.), *Handbook of research methods in social and personality psychology* (2nd edition, pp. 283–310). New York: Cambridge University Press.
- Gawronski, B., Morrison, M., Phillips, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, *43*, 300–312.
- Gawronski, B., & Sritharan, R. (2010). Formation, change, and contextualization of mental associations: Determinants and principles of variations in implicit measures. In B. Gawronski, & B. K. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 216–240). New York: Guilford Press.
- Greenwald, A. G., Poehlman, T. A., Uhlmann, E., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97*, 17–41.
- Gschwendner, T., Hofmann, W., & Schmitt, M. (2008). Differential stability: The effects of acute and chronic construct accessibility on the temporal stability of the Implicit Association Test. *Journal of Individual Differences*, *29*, 70–79.
- Koch, T., Ortner, T. M., Eid, M., Caspers, J., & Schmitt, M. (2014). Evaluating the construct validity of objective personality tests using a multitrait-multimethod-multioccasion-(MTMM-MO)-approach. *European Journal of Psychological Assessment*, *30*, 208–230.
- Kurdi, B., & Dunham, Y. (2020). Propositional accounts of implicit evaluation: Taking stock and looking ahead. *Social Cognition*, *38*, s42–s67.
- Kurdi, B., Seitchik, A. E., Axt, J. R., Carroll, T. J., Karapetyan, A., Kaushik, N., Tomezsko, D., Greenwald, A. G., & Banaji, M. R. (2019). Relationship between the Implicit Association Test and intergroup behavior: A meta-analysis. *American Psychologist*, *74*, 569–586.
- Lemmer, G., Gollwitzer, M., & Banse, R. (2015). On the psychometric properties of the aggressiveness-IAT for children and adolescents. *Aggressive Behavior*, *41*, 84–95.
- Machery, E. (2021). Anomalies in implicit attitude research. *Wiley Interdisciplinary Reviews: Cognitive Science*. Advance online publication. <https://doi.org/10.1002/wcs.1569>
- Madva, A. (2020). Resistance training. *The Philosophers' Magazine*, *91*, 40–45.
- Mallon, R. (2021). Racial attitudes, accumulation mechanisms, and disparities. *Review of Philosophy and Psychology*. Advance online publication. <https://doi.org/10.1007/s13164-020-00521-6>
- Oswald, F. L., Mitchell, G., Blanton, H., Jaccard, J., & Tetlock, P. E. (2013). Predicting ethnic and racial discrimination: A meta-analysis of IAT criterion studies. *Journal of Personality and Social Psychology*, *105*, 171–192.
- Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2017). The bias of crowds: How implicit bias bridges personal and systemic prejudice. *Psychological Inquiry*, *28*, 233–248.
- Schmukle, S. C. & Egloff, B. (2005). A latent state-trait analysis of implicit and explicit personality measures. *European Journal of Psychological Assessment*, *21*, 100–107.
- Schwarz, N. (2007). Attitude construction: Evaluation in context. *Social Cognition*, *25*, 638–656.
- Sherman, J. W., Klauer, K. C., & Allen, T. J. (2010). Mathematical modeling of implicit social cognition: The machine in the ghost. In B. Gawronski & B. K. Payne (Eds.), *Handbook of implicit social cognition: Measurement, theory, and applications* (pp. 156–175). New York: Guilford Press.
- Spencer, S. J., Zanna, M. P., & Fong, G. T. (2005). Establishing a causal chain: Why experiments are often more effective than mediational analyses in examining psychological processes. *Journal of Personality and Social Psychology*, *89*, 845–851.
- Strack, F., & Deutsch, R. (2004). Reflective and impulsive determinants of social behavior. *Personality and Social Psychology Review*, *8*, 220–247.