PSYCHOLOGICAL INQUIRY
2017, VOL. 28, NO. 4, 268–272
https://doi.org/10.1080/1047840X.2017.1373546

Routledge
Taylor & Francis Group

Beyond Persons and Situations: An Interactionist Approach to Understanding Implicit Bias

Bertram Gawronski[a] and Galen V. Bodenhausen[b]

[a]Department of Psychology, University of Texas at Austin, Austin, Texas; [b]Department of Psychology, Northwestern University, Evanston, Illinois

The dominant view in the public and scientific discourse about implicit bias has shifted remarkably in the last few years. Counter to earlier views that identified expressions of implicit bias in almost every instance of social discrimination, many researchers have become rather skeptical of the construct—up to a point where it is dismissed by some as being entirely irrelevant for the psychological understanding of discriminatory behavior. This growing skepticism, which affected even some of the early proponents of the construct (e.g., Forscher, Mitamura, Dix, Cox, & Devine, in press), has its roots in at least two sets of empirical findings. First, implicit biases are relatively unstable, showing considerable fluctuations over short periods of time (e.g., Cunningham, Preacher, & Banaji, 2001; Gawronski, Morrison, Phills, & Galdi, 2017). Second, relations between measures of implicit bias and individual differences in discriminatory behavior tend to be relatively weak, with meta-analytic correlations ranging from $r = .14$ to .28 (Cameron, Brown-Iannuzzi, & Payne, 2012; Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013). As noted by Payne, Vuletich, and Lundberg (this issue), these findings seem at odds with other evidence suggesting that (a) implicit biases are strong and widespread on average (e.g., Nosek et al., 2007; Payne, Krosnick, Pasek, Lelkes, Akhtar & Tompson, 2010), (b) children of very young age show levels of implicit bias that are indistinguishable from the levels revealed by adults (e.g., Banse, Gawronski, Rebetez, Gutt, & Morton, 2010; Baron & Banaji, 2006), and (c) aggregate scores of implicit bias at the macrolevel (e.g., regions, states, countries) show strong relations with indicators of discrimination at the same level of analysis (e.g., Leitner, Hehman, Ayduk, & Mendoza-Denton, 2016; Nosek et al., 2009).

Payne et al.'s (this issue) Bias of Crowds (BoC) model resolves the apparent conflicts between these findings by treating implicit biases as reflections of situation-related rather than person-related factors. The basic idea is that implicit biases are the product of social disparities at the macrolevel, which lead to context-dependent fluctuations in the accessibility of mental concepts at the individual level. According to this view, implicit biases reflect the momentary accessibility of mental concepts associated with different groups, which is shaped by the level of prejudice and discrimination in a person's environment. By attributing a dominant role to situational (in contrast to personal) determinants of concept accessibility (see Higgins, 1996), the BoC model resolves the apparent conflicts between the preceding sets of findings. Issues of cognitive architecture are deemed irrelevant, in that the BoC model is compatible with a broad range of theoretical assumptions about mental representation.

The BoC model provides a sophisticated and refreshingly balanced way to think about implicit bias, which could help to reconcile the polarized views in the current debate. However, despite our agreement with many of the model's core assumptions, we are concerned about its "blind spot" regarding central issues of cognitive architecture. Our main argument is that, by ignoring the mental processes and representations that determine concept accessibility, the BoC model misses the significance of person-related factors in the activation of mental contents. In our view, concept accessibility must be understood as the interactive product of person-related and situation-related factors, which leads to a more nuanced understanding of implicit biases, their determinants, and their behavioral consequences.

Person × Situation Interactions

Different from the BoC model and other conceptions that treat personal and situational factors as additive determinants of accessibility (e.g., Bargh, Bond, Lombardi, & Tota, 1986), we argue that accessibility has to be understood as the interactive product of (a) the preexisting structure of associations in memory (person-related factor) and (b) the overall configuration of input stimuli (situation-related factor). Within the framework of our associative-propositional evaluation (APE) model (Gawronski & Bodenhausen, 2006), we proposed the following:

> Association activation is not an all-or-none process, such that encountering a given object activates each and every mental association related to that object. Instead, objects tend to activate only a limited subset of all object-related associations that are available in memory. Which subset of associations is activated in response to a given object is assumed to depend on the overall configuration of input stimuli. For example, encountering a Black person in a jazz bar may activate the stereotypical attribute *musical*, whereas the same Black person may activate the stereotypical attribute *criminal* if that person is encountered in a dark alley. (p. 62)

This idea captures the tenet of the BoC model that context plays a central role for the activation of mental contents in

CONTACT Bertram Gawronski gawronski@utexas.edu Department of Psychology, University of Texas at Austin, 108 E Dean Keeton A8000, Austin, TX 78712, USA.

response to a given stimulus. Yet, the APE model goes the BoC model by assuming that the overall configuration of input stimuli interacts with a person's mental associations in determining the activation of mental contents. In the preceding example, encountering a Black person in a jazz bar would activate the concept *musical* only if the perceiver's mental representation includes an association between *African Americans* and *musical*. Similarly, encountering a Black person in a dark alley would activate the concept *criminal* only if the perceiver's mental representation includes an association between *African Americans* and *criminal*. Thus, different from the assumption that personal and situational factors influence concept accessibility in an additive manner, we propose that activation of mental contents must be understood as the product of Person × Situation interactions. From an interactionist view, situational effects on concept accessibility still involve a person-related component, in that they operate on preexisting associations in memory. To the extent that the structure and relative strength of mental associations varies across individuals, given each person's idiosyncratic learning history, the same situational input can activate different concepts for different individuals. At the same time, concept accessibility is not a context-independent product of preexisting associations, in that the same stimulus can activate different mental contents depending on the context in which it is encountered. Thus, we agree with the BoC model that a narrow focus on person-related factors provides an insufficient understanding of implicit bias. We also agree that a sufficient understanding of implicit biases requires a strong emphasis of situational factors. However, we disagree with the BoC model's assumption that implicit bias can be understood as the exclusive product of situation-related factors without any consideration of person-related factors. Our main argument is that implicit bias must be understood as the joint product of person-related and situation-related factors, involving interactive effects of preexisting associations and overall configurations of input stimuli in the activation of mental contents.

The significance of these issues can be illustrated with a largely neglected study by Gschwendner, Hofmann, and Schmitt (2008). Consistent with many other studies (e.g., Cunningham et al., 2001; Gawronski et al., 2017), the authors found rather low levels of stability in implicit bias over a period of two weeks when they used a standard variant of the Implicit Association Test ($r = .29$). Yet, temporal stability of implicit bias over the same period was significantly higher when the measure included background images to provide additional information about the context of the target stimuli ($r = .72$).[1] These findings suggest that a person's level of implicit bias fluctuates over time in the absence of strong contextual constraints. However, implicit bias seems to be quite stable over time to the extent that contextual constraints are strong and consistent across measurements.

These findings have important implications for the BoC model's hypothesis that concept accessibility is the exclusively product of situation-related factors with little impact of person-related factors. If context was the only source of systematic variance in implicit bias and person-related factors did not matter, participants in Gschwendner et al. (2008) study should show similar scores of implicit bias when the context of the target stimuli was held constant across participants. Any variation in implicit bias scores across participants would be measurement error, because there should be no meaningful individual differences. Hence, correlations between implicit bias scores obtained at different time points should be lower in the presence (vs. absence) of contextual constraints because there are no systematic differences in implicit bias across participants. Yet, Gschwendner et al. found exactly the opposite. When context was specified and held constant across participants, individual scores of implicit bias showed higher (not lower) correlations over time. This result suggests that (a) a given stimulus can activate different mental contents for different people even when the stimulus is encountered in the same context, and (b) individual differences in the effect of stimulus-context configurations are relatively stable over time.

To avoid potential confusion, it is worth clarifying which particular assumptions of the BoC model are challenged by Gschwendner et al. (2008) findings and which ones are not. In our view, the BoC model is correct in its assumption that there is considerable situational variability in concept accessibility at the individual level, which leads to low stability of implicit bias over time (because contexts tend to change over time). The BoC is also correct in its assumption that the high stability of implicit bias at the aggregate level (e.g., regions, states, countries) is best explained by a strong and relatively persistent impact of context. However, we argue that the BoC model prematurely dismisses the role of person-related factors by ignoring the role of Person × Situation interactions. Although situations can have powerful effects on people, people tend to differ in their response to the same situation, and individual differences in situation-behavior profiles can be relatively stable over time (Mischel & Shoda, 1995). Applied to the current question, the latter assumption is consistent with the finding that implicit bias is more stable over time when the context is taken into account (Gschwendner et al., 2008). It is also consistent with weak predictive relations between implicit bias and discriminatory behavior, which can be due to changes in the context (i.e., when implicit bias and discriminatory behavior are measured in different contexts). Finally, it is consistent with the finding that aggregate scores of implicit bias tend to be relatively stable at the macrolevel (i.e., when implicit bias scores are analyzed for different age or demographic groups). After all, mean scores of implicit bias can be relatively stable at the aggregate level even when there is temporal and contextual variation at the individual level (for a discussion, see Gawronski et al., 2017).

Change in Association Activation versus Change in Associative Structure

Our arguments about the interactive role of input configurations and preexisting associations also have important implications for observed changes in a person's level of implicit bias as a result of situational influences. According to our APE model, such changes may reflect a change in either (a) the momentary activation of existing associations in response to a given object

[1] Similar findings were obtained for an Implicit Association Test designed to measure the implicit self-concept of anxiety.

or (b) the structure or strength of mental associations related to that object (Gawronski & Bodenhausen, 2006, 2011). The former can be described as a *context effect* in the sense that the same target object may activate different patterns of preexisting associations depending on the context in which it is encountered. The latter may be described as a *learning effect* in the sense that it involves a genuine change in the mental representation of the target object. Although the two are often lumped together in research on change in implicit bias, we deem it important to distinguish between the two instances of "change" because context effects and learning effects have fundamentally different properties. For example, although temporally stable changes in implicit bias may be expected when such changes are driven by a change in the structure or strength of mental associations in memory (e.g., Olson & Fazio, 2006), there is no basis to expect temporally stable changes in implicit bias when the observed changes result from different contextual constraints on the activation of preexisting associations (e.g., Ma, Correll, & Wittenbrink, 2016).

Although a change in a person's associative network seems necessary for temporally stable changes in implicit bias, it is not sufficient. First, newly formed associations would have to be strong enough to have long-lasting effects. To the extent that a newly formed association is relatively weak, its contribution to concept accessibility at future encounters with the target object may be relatively small (Petty & Cacioppo, 1986). Second, even when a newly formed association is strong, its impact on implicit bias can be limited to the context in which the association was formed. Thus, to the extent that the target object is encountered in a context that is different from the one in which the new association was formed, its contribution to concept accessibility may be limited even when this association is relatively strong (Gawronski & Cesario, 2013).

The latter idea is reflected in research on contextual renewal effects in the formation and change of implicit evaluations (e.g., Gawronski, Rydell, Vervliet, & De Houwer, 2010; Rydell & Gawronski, 2009; Ye, Tong, Chiu, & Gawronski, 2017; for a review, see Gawronski et al., in press). Borrowed from research on animal learning, *contextual renewal* describes the return of a previously acquired response to a given stimulus (e.g., a conditioned fear response to a tone) when the stimulus is encountered in a context that differs from the one in which a new response to that stimulus was learned (e.g., after extinction or counterconditioning of the fear response; for reviews, see Bouton, 2004; Gawronski & Cesario, 2013; Vervliet, Baeyens, Van den Bergh, & Hermans, 2013a). Applied to changes in implicit bias, contextual renewal would occur if a bias intervention effectively reduced implicit bias in the context in which the intervention occurred but not in contexts that are different from the intervention context.

Research on contextual renewal has important implications for implicit bias because it provides a potential explanation for why many bias interventions tend to be ineffective in producing persistent changes that are stable over time (e.g., Lai et al., 2016). To the extent that follow-up assessments occur in contexts that are different from the one in which the intervention occurred, a recurrence of implicit bias may not necessarily reflect ineffectiveness of the intervention in producing long-term change. After all, the recurrence of implicit bias could also be due to the change in context, such that the intervention may be effective in producing long-term change in implicit bias within the context in which the intervention occurred. In other words, the problem of ineffective interventions may not necessarily be a low temporal stability of the observed change. Instead, the problem may be a high level of context-specificity in the observed change. That is, the observed change may be highly stable within the context in which the intervention occurred, but this change may not generalize across contexts (see Vervliet, Craske, & Hermans, 2013b). Because lack of temporal stability and lack of generalization result from different aspects of newly formed associations (i.e., weak association vs. contextualized association), they also require different modifications of existing bias interventions to increase their overall effectiveness.

Implicit Bias and Discriminatory Behavior

An interactionist view also has important implications for understanding the weak relation between implicit bias and individual differences in discriminatory behavior (for meta-analyses, see Cameron et al., 2012; Greenwald et al., 2009; Oswald et al., 2013). To the extent that implicit bias and the to-be-predicted behavior are measured in different contexts, weak predictive relations are not particularly surprising. Such changes in context are particularly likely when there is a delay in the measurement of implicit bias and discriminatory behavior. In such cases, weak predictive relations can be explained by the same mechanisms that account for the low stability of implicit bias over time. Because contexts tend to change over time and activation of mental contents can differ across contexts, implicit bias and discriminatory behavior may show relatively weak relations when they are measured at different time points. However, this does not mean that a person's level of implicit bias is unrelated to discriminatory behavior when they are measured within the same context (regardless of potential delays).

Yet, even in such cases, it seems unlikely that implicit bias shows unconditional relations to discriminatory behavior. After all, predictive relations between implicit bias and discriminatory behavior have been found to depend on the type of behavior, the conditions under which the behavior is performed, and characteristics of the person who is performing the behavior (for a review, see Friese, Hofmann, & Schmitt, 2008). In the broader literature on implicit measures, these boundary conditions are reflected in studies showing that (a) implicit measures tend to show stronger relations to spontaneous compared to deliberate behavior (e.g., Fazio, Jackson, Dunton, & Williams, 1995), (b) implicit measures tend to show stronger relations to behavior when the behavior is performed under conditions of low compared to high cognitive capacity (e.g., Hofmann, Gschwendner, Friese, Wiers, & Schmitt, 2008), and (c) implicit measures tend to show stronger relations to behavior for individuals with a preference for intuitive compared to deliberate thinking styles (e.g., Richetin, Perugini, Adjali, & Hurling, 2007). Explicit measures have been found to show the opposite patterns in the prediction of behavior (for a review, see Friese et al., 2008). Thus, when relations between implicit bias and

discriminatory behavior are analyzed at the level of zero-order correlations ignoring the aforementioned moderators (e.g., type of behavior, cognitive resources, individual differences in thinking styles), the to-be-expected outcome is a weak positive relation, as found in extant meta-analyses (e.g., Cameron et al., 2012; Greenwald et al., 2009; Oswald et al., 2013). These findings conceal the fact that relations between implicit bias and discriminatory behavior at the individual level tend to be quite strong for particular behaviors, particular conditions, and particular kinds of people, and relatively weak for other behaviors, other conditions, and other kinds of people.[2] These issues go beyond the main arguments of the BoC model, which emphasizes strong relations between aggregate scores of implicit bias and social discrimination at the macrolevel (i.e., regions, states, countries).

Conclusion

Conceptually, the basic idea of the BoC model involves an aggregation of measurement scores across individuals to identify effects of situation-related factors (see Surowiecki, 2004). This idea resembles earlier approaches in research on attitudes (e.g., Ajzen, 1987) and personality (e.g., Epstein, 1979), involving aggregations of measurement scores across situations to identify effects of person-related factors. In our view, either of the two approaches provides a limited understanding of human behavior. Aggregation across individuals to identify effects of situation-related factors ignores the significance of person-related factors; aggregation across situations to identify effects of person-related factors ignores the significance of situation-related factors. Thus, neither of the two approaches is able to capture the significance of Person × Situation interactions (Mischel & Shoda, 1995), which we deem essential for a sufficient understanding of implicit bias. Different people show different responses to a given stimulus within same situation, and the same person may show different responses to the same stimulus depending on the context in which the stimulus is encountered. By treating concept accessibility as an interactive product of person-related and situation-related factors, an interactionist view not only integrates the valuable insights provided by the BoC model but also captures various other findings that remain unexplained by the BoC model (e.g., high temporal stability of implicit bias within the same context; contextual renewal effects in the formation and change of implicit bias; conditional nature of predictive relations at the individual level). This capacity emerged from the focus on mental processes and representations underlying concept accessibility, which provides more a fine-grained understanding of dynamic interactions between person-related and situation-related factors. Implicit bias is not a simple characteristic of the person, and it is not a pure reflection the situation; it is a reflection of the person *within* a given situation.

[2]To our surprise, these issues have been largely ignored in the current debate about implicit bias, although the conditional nature of attitude-behavior relations has been extensively discussed in 1970s and 1980s (e.g., Zanna & Fazio, 1982). Similar to the current argument, the main conclusion from these earlier debates was that the strength attitude-behavior relations depend on the particular behavior, the conditions under which the behavior is performed, and the person who is performing the behavior.

References

Ajzen, I. (1987). Attitudes, traits, and actions: Dispositional prediction of behavior in personality and social psychology. *Advances in Experimental Social Psychology*, *20*, 1–63.

Banse, R., Gawronski, B., Rebetez, C., Gutt, H., & Morton, J. B. (2010). The development of spontaneous gender stereotyping in childhood: Relations to stereotype knowledge and stereotype flexibility. *Developmental Science*, *13*, 298–306. doi:10.1111/j.1467-7687.2009.00880.x

Bargh, J. A., Bond, R. N., Lombardi, W. J., & Tota, M. E. (1986). The additive nature of chronic and temporary sources of construct accessibility. *Journal of Personality and Social Psychology*, *50*, 869–878. doi:10.1037/0022-3514.50.5.869

Baron, A. S., & Banaji, M. R. (2006). The development of implicit attitudes: Evidence of race evaluations from ages 6 and 10 and adulthood. *Psychological Science*, *17*, 53–58. doi:10.1111/j.1467-9280.2005.01664.x

Bouton, M. E. (2004). Context and behavioral processes in extinction. *Learning and Memory*, *11*, 485–494. doi:10.1101/lm.78804

Cameron, C. D., Brown-Iannuzzi, J. L., & Payne, B. K. (2012). Sequential priming measures of implicit social cognition: A meta-analysis of associations with behavior and explicit attitudes. *Personality and Social Psychology Review*, *16*, 330–350. doi:10.1177/1088868312440047

Cunningham, W. A., Preacher, K. J., & Banaji, M. R. (2001). Implicit attitude measurement: Consistency, stability, and convergent validity. *Psychological Science*, *12*, 163–170. doi:10.1111/1467-9280.00328

Epstein, S. (1979). The stability of behavior: I. On predicting most of the people most of the time. *Journal of Personality and Social Psychology*, *37*, 1097–1126. doi:10.1037/0022-3514.37.7.1097

Fazio, R. H., Jackson, J. R., Dunton, B. C., & Williams, C. J. (1995). Variability in automatic activation as an unobtrusive measure of racial attitudes: A bona fide pipeline? *Journal of Personality and Social Psychology*, *69*, 1013–1027. doi:10.1037/0022-3514.69.6.1013

Forscher, P. S., Mitamura, C., Dix, E. L., Cox, W. T. L., & Devine, P. G. (in press). Breaking the prejudice habit: Mechanisms, time course, and longevity. *Journal of Experimental Social Psychology*.

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.

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. doi:10.1037/0033-2909.132.5.692

Gawronski, B., & Bodenhausen, G. V. (2011). The associative-propositional evaluation model: Theory, evidence, and open questions. *Advances in Experimental Social Psychology*, *44*, 59–127. doi:10.1016/B978-0-12-385522-0.00002-0

Gawronski, B., & Cesario, J. (2013). Of mice and men: What animal research can tell us about context effects on automatic responses in humans. *Personality and Social Psychology Review*, *17*, 187–215. doi:10.1177/1088868313480096

Gawronski, B., Morrison, M., Phills, C. E., & Galdi, S. (2017). Temporal stability of implicit and explicit measures: A longitudinal analysis. *Personality and Social Psychology Bulletin*, *43*, 300–312. doi:10.1177/0146167216684131

Gawronski, B., Rydell, R. J., De Houwer, J., Brannon, S. M., Ye, Y., Vervliet, B., & Hu, X. (in press). Contextualized attitude change. *Advances in Experimental Social Psychology*.

Gawronski, B., Rydell, R. J., Vervliet, B., & De Houwer, J. (2010). Generalization versus contextualization in automatic evaluation. *Journal of Experimental Psychology: General*, *139*, 683–701. doi:10.1037/a0020315

Greenwald, A. G., Poehlman, T. A., Uhlmann, E. L., & Banaji, M. R. (2009). Understanding and using the Implicit Association Test: III. Meta-analysis of predictive validity. *Journal of Personality and Social Psychology*, *97*, 17–41. doi:10.1037/a0015575

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. doi:10.1027/1614-0001.29.2.70

Higgins, E. T. (1996). Knowledge activation: Accessibility, applicability, and salience. In E. T. Higgins & A. W. Kruglanski (Eds.), *Social psychology: Handbook of basic principles* (pp. 133–168). New York, NY: Guilford Press.

Hofmann, W., Gschwendner, T., Friese, M., Wiers, R., & Schmitt, M. (2008). Working memory capacity and self-regulatory behavior: Towards and individual differences perspective on behavior determination by automatic versus controlled processes. *Journal of Personality and Social Psychology*, *95*, 962–977. doi:10.1037/a0012705

Lai, C. K., Skinner, A. L., Cooley, E., Murrar, S., Brauer, M., Devos, T., … Nosek, B. A. (2016). Reducing implicit racial preferences: II. Intervention effectiveness across time. *Journal of Experimental Psychology: General*, *145*, 1001–1016.

Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016). Racial bias is associated with ingroup death rate for Blacks and Whites: Insights from Project Implicit. *Social Science & Medicine*, *170*, 220–227. doi:10.1016/j.socscimed.2016.10.007

Ma, D. S., Correll, J., & Wittenbrink, B. (2016). Context dependency at recall: Decoupling context and targets at encoding. *Social Cognition*, *34*, 119–132. doi:10.1521/soco.2016.34.2.119

Mischel, W., & Shoda, Y. (1995). A cognitive-affective system theory of personality: Reconceptualizing situations, dispositions, dynamics, and invariance in personality structure. *Psychological Review*, *102*, 246–268. doi:10.1037/0033-295X.102.2.246

Nosek, B. A., Smyth, F. L, Hansen, J. J., Devos, T., Lindner, N. M., Ranganath, K. A., … Banaji, M. R. (2007). Pervasiveness and correlates of implicit attitudes and stereotypes. *European Review of Social Psychology*, *18*, 36–88. doi:10.1080/10463280701489053

Nosek, B. A., Smyth, F. L., Sriram, N., Linder, N. M., Devos, T., Ayala, A., … Greenwald, A. G. (2009). National differences in gender-science stereotypes predict national sex differences in science and math achievement. *Proceedings of the National Academy of Sciences*, *106*, 10593–10597. doi:10.1073/pnas.0809921106

Olson, M. A., & Fazio, R. H. (2006). Reducing automatically activated racial prejudice through implicit evaluative conditioning. *Personality and Social Psychology Bulletin*, *32*, 421–433. doi:10.1177/0146167205284004

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. doi:10.1037/a0032734

Payne, B. K., Krosnick, J. A., Pasek, J., Lelkes, Y., Akhtar, O., & Tompson, T. (2010). Implicit and explicit prejudice in the 2008 American presidential election. *Journal of Experimental Social Psychology*, *46*, 367–374. doi:10.1016/j.jesp.2009.11.001

Petty, R. E., & Cacioppo, J. T. (1986). The elaboration likelihood model of persuasion. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 19, pp. 123–205). New York, NY: Academic Press.

Richetin, J., Perugini, M., Adjali, I., & Hurling, R. (2007). The moderator role of intuitive versus deliberative decision making for the predictive validity of implicit measures. *European Journal of Personality*, *21*, 529–546. doi:10.1002/per.625

Rydell, R. J., & Gawronski, B. (2009). I like you, I like you not: Understanding the formation of context-dependent automatic attitudes. *Cognition and Emotion*, *23*, 1118–1152. doi:10.1080/02699930802355255

Surowiecki, J. (2004). *The wisdom of crowds: Why the many are smarter than the few and how collective wisdom shapes business, economies, societies, and nations*. New York, NY: Doubleday Books.

Vervliet, B., Baeyens, F., Van den Bergh, O., & Hermans, D. (2013a). Extinction, generalization, and return of fear: A critical review of renewal research in humans. *Biological Psychology*, *92*, 51–58. doi:10.1016/j.biopsycho.2012.01.006

Vervliet, B., Craske, M. G., & Hermans, D. (2013b). Fear extinction and relapse: State of the art. *Annual Review of Clinical Psychology*, *9*, 215–248.

Ye, Y., Tong, Y.-Y., Chiu, C.-Y., & Gawronski, B. (2017). Attention to context during evaluative learning and context-dependent automatic evaluation: A cross-cultural analysis. *Journal of Experimental Social Psychology*, *70*, 1–7.

Zanna, M. P., & Fazio, R. H. (1982). The attitude-behavior relation: Moving toward a third generation of research. In M. P. Zanna, E. T. Higgins, & C. P. Herman (Eds.), *Consistency in social behavior: The Ontario Symposium* (Vol. 2, pp. 283–301). Hillsdale, NJ: Erlbaum.