UC Davis
UC Davis Previously Published Works

Title
Implicit Bias Reflects the Personal and the Social

Permalink
https://escholarship.org/uc/item/1s32872b

Authors
Rivers, Andrew M
Rees, Heather
Calanchini, Jimmy
et al.

Publication Date
2018-01-19

DOI
10.31234/osf.io/qm592

Peer reviewed
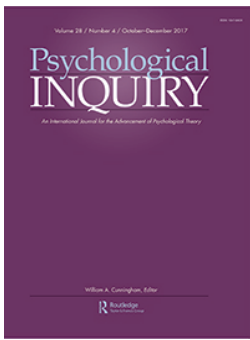
Implicit Bias Reflects the Personal and the Social

Andrew M. Rivers, Heather R. Rees, Jimmy Calanchini & Jeffrey W. Sherman

Published online: 30 Nov 2017.

Submit your article to this journal ↗

View related articles ↗

View Crossmark data ↗

Implicit Bias Reflects the Personal and the Social

Andrew M. Rivers[a], Heather R. Rees[a], Jimmy Calanchini[b], and Jeffrey W. Sherman[a]

[a]Department of Psychology, University of California, Davis, Davis, California; [b]Institut für Psychologie, Albert-Ludwigs-Universität Freiburg, Freiburg im Breisgau, Germany

This issue's target article by Payne, Vuletich, and Lundberg (PV&L) does exactly what one should, presenting an argument that is thought-provoking and that challenges current orthodoxy. It also addresses an issue that has increasingly confounded attitudes researchers in recent years.

The construct of "implicit bias" was initially conceptualized as a latent construct that exists within persons, relatively resistant to situational influences. A plethora of theoretical models converge on the notion that implicit biases, including intergroup biases, are representations that are stored in memory (e.g., Devine, 1989; Fazio, Jackson, Dunton, & Williams, 1995; Gawronski & Bodenhausen, 2006; Greenwald et al., 2002; Wilson, Lindsay, & Schooler, 2000). Although some perspectives emphasize the role of culture in contributing to implicit measures of bias, even these perspectives rely on the learning and storage of mental representations (Olson & Fazio, 2004).

Researchers have sought to bridge the divide separating biases of individual perceivers from the larger culture within which they are embedded (e.g., Dunham, Chen, & Banaji, 2013). This has been no easy task, which may have been stymied by the early conceptualization of implicit bias as a stable individual difference, rather than as a construct prone to situational malleability. PV&L (this issue) offer Bias of Crowds (BoC), an intriguing perspective on implicit bias that provocatively pushes situational variability as central and primary, rather than secondary. We applaud this new perspective; after all, we are social psychologists with a long tradition of studying how situations influence human cognition and behavior. Although the BoC emphasis on situational influence offers a much-needed shift toward the social component of implicit bias, we believe that the strong claim that implicit measures reflect only social context goes too far, and is not supported by the data.

Individual versus Aggregate Level Implicit Bias and Behavior

We share PV&L's enthusiasm for the aggregate level (e.g., nation, state, metropolitan area) approach to examining the relationship between implicit bias and behavioral outcomes. Indeed, some of us are actively involved in such investigations (e.g., Calanchini, Hehman, & Salerno, 2017; Calanchini, Witkowski, Sparks, Schmidt, & Sherman, 2017; Hehman, Calanchini, Flake & Leitner, 2017; Hehman, Flake, & Calanchini, in

press). However, our willingness to draw strong conclusions about the superiority of implicit bias operationalized at the aggregate versus individual level in predicting behavior is tempered by the relatively small number of studies to date that have employed the aggregate approach. To our knowledge, PV&L's summary of published research on aggregate implicit bias is exhaustive, with the exception of one article by Leitner, Hehman, Ayduk, and Mendoza-Denton (2016a), who found that county level implicit racial bias predicted regional racial mortality disparities, but only when implicit and explicit biases were modeled separately. Whereas only a handful of scientific publications have so far examined the relationship between aggregate implicit bias and behavior, hundreds of studies have examined the relationship between individual implicit bias and behavior, which have been summarized in multiple meta-analyses (e.g., Cameron, Brown-Iannuzzi, & Payne, 2012; Greenwald, Poehlman, Uhlmann, & Banaji, 2009; Oswald, Mitchell, Blanton, Jaccard, & Tetlock, 2013). Consequently, we have a very good idea of how strongly individual implicit bias corresponds to behavior, with estimates ranging from $r = .14$ (Oswald et al., 2013) to .28 (Cameron et al., 2012). However, given that the upper limit on which two measures can correlate is determined by the reliability of each measure, we cannot unambiguously interpret these low correlations between individual implicit bias and behavior as evidence of either low attitude/behavior correspondence or poor measurement reliability. Region level analyses have thus far produced relatively large effects and tend to be very well powered (e.g., Hehman et al., 2017; Leitner et al., 2016a, 2016b; Nosek et al., 2009; Zerhouni, Rougier, & Muller, 2016). Because aggregation increases reliability, these initial region level findings suggest that individual level attitude/behavior correlations may be suppressed by low measurement reliability. Nevertheless, more evidence is needed before we can draw strong conclusions about the superiority of aggregate- versus individual level predictions of biased behavior.

BoC versus Models of Chronic and Temporary Accessibility

We believe that BoC could be more effectively integrated with existing models of chronic and temporary accessibility. As PV&L acknowledge, there is a long history of research examining conditions under which implicit intergroup biases are

situationally malleable (e.g., Blair, 2002). In addition, many existing models of accessibility allow for the influence of both stable stored target-attitude relations and temporarily accessible or activated relations. More broadly, social psychologists have studied the interactions between chronically and temporarily accessible information for a long time. For example, Higgins (1996) proposed that social contexts increase the temporary accessibility of context-relevant information, whereas chronic accessibility is rooted in the individual's learning history, including cultural contexts, personally endorsed lay theories, and the frequency of exposure to the concept. Similarly, Fazio (2007) asserted that attitudes can be thought of as object-evaluation associations that are somewhat stable and differ in strength based on prior learning history but that the extent to which these attitudes are *expressed* in behavior depends both on the strength of the stable association and on what features are made contextually salient. Conrey and Smith's (2007) distributed connectionist model of attitudes expresses temporary and chronic forms of accessibility in a plausible self-contained system and illustrates how repeated accessibility can eventually lead to change in overlearned, chronically accessible attitudes. Each of these models acknowledges that situations can temporarily influence the accessibility of specific concepts but also balances this against the individual's learning history. Indeed, few researchers would disagree with the critical assumption of PV&L's proposal: Situations meaningfully influence implicit biases at the individual level. When considered in comparison to existing attitude models, it is clear that the novelty of BoC is not so much in its claims of situational dependence but, instead, in the weight it assigns to situations relative to individual learning histories.

Just as BoC largely fits into the existing constellation of attitude models, so too can many of the puzzles highlighted by PV&L be understood through the lens of standard accessibility models. For example, PV&L note that the question of whether implicit bias is personal or extrapersonal is rendered moot when implicit biases are exclusively the product of temporary accessibility. However, models that permit the influence of both chronically and temporarily accessible information find the whole binary nature of the proposal simplistic, and not in need of resolution in the first place. First, at any given time, bias will necessarily reflect a combination of what is chronic and what is temporary. Second, there is no reason to doubt that individual implicit biases reflect, among other things, the extent to which culture has codified inequality into the societal fabric (e.g., laws and stable cultural norms). The either/or demand that implicit bias reflect either the personal or the social is strange. No one questions whether Americans' general aversion to horse meat is truly an individual attitude versus a cultural construct. Culture affects the development of persistent associations as well as what concepts are temporarily accessible. This fact does not change the conclusion that biases are, in part, representations residing within individuals.

Likewise, explaining the (un)consciousness of implicit biases does not require the extreme contextualist proposal of BoC. The fact that people can accurately predict their implicit biases requires only that they can recognize what concepts are currently accessible or, alternatively, recognize the spontaneous biases that have been activated. Accessibility could arise from either temporary sources such as recently activated concepts or chronic attitudes that are rooted in learning history (or more likely a mixture of both). The research discussed by PV&L (Hahn et al., 2014; Ranganath, Smith, & Nosek, 2008) is unable to distinguish between these two sources. Similarly, that people report different biases when asked for considered opinions or instant reactions does not speak to whether implicit biases are determined by temporary or chronic representations. Instead, the existing research indicates only that people recognize that their spontaneous reactions may differ from what they personally endorse and/or what they report on explicit measures.

A situational approach to interventions that target implicit biases is well worth considering but again is not incompatible with existing accessibility models. PV&L suggest that the situationist approach to implicit bias interventions needs more research. We wholeheartedly agree with this recommendation and also agree that, thus far, interventions focusing on changing individual level implicit intergroup biases have been ineffective (e.g., Lai et al., 2016). A revised accessibility model that assigns substantial weight to temporary influences (and little to chronic factors) would converge with BoC in targeting situations, more so than individuals, for intervention. To the extent that temporary (versus chronic) factors dictate concept accessibility, interventions that target the causes of temporary accessibility will be most fruitful. As PV&L note, there appear to be attitudinal domains (e.g., political attitudes) in which chronic accessibility may play a larger role compared with intergroup attitudes. These cases can also be understood through the lens of a traditional accessibility model that assigns a greater degree of weight to chronic influences and do not necessitate a model giving all the weight to current context. In general, we see no advantage to developing different attitude models to account for attitudes in different domains.

Where Did the Bias Come From?

One challenge for BoC is accounting for the initial formation of bias. BoC can explain how implicit intergroup biases are perpetuated once established, but seems to have challenges explaining the ultimate causes for intergroup biases. In other words, how do evaluative biases emerge in the first place? It is clear that biased crowds may codify their intergroup preferences into regulations, norms, or formal lawmaking. This would serve as a stable situational cue to influence subsequent perceivers. But at some point in the development of intergroup preferences, an individual or group must have learned intergroup bias without having been taught. Further, the individual's preference must have been durable enough to persist over time and communicable enough to spread widely within that individual's social circle.

There are many accounts for the development of intergroup preferences, including evolutionary accounts, learning and attentional processes, and intergroup competition, among others. Most, if not all, of these accounts would agree with the idea that social learning serves as a mechanism to propagate biases across cultures but would disagree with the idea that individuals do not store stable intergroup attitudes. PV&L imply that implicit biases flow primarily from stable systems, in which bias is formalized, to the individual who is engaged with

the biased system. But, at some point, individuals must have successfully influenced the system to formally instantiate its biases. If not, another mechanism is necessary to explain what led to the initial formalization of intergroup biases. Models that allow for some attitude stability within individuals are better equipped to explain the early formalization of systematic biases.

Accessibility and Control at Individual and Aggregate Levels of Analysis

Although situations surely influence the accessibility of different concepts and associations, they also influence measured implicit bias via different processes. Our own research, as well as that of PV&L's research group and others, has demonstrated that implicit measures reflect more than accessible associations (e.g., associations between Black–unpleasant, White–pleasant). Instead, responses on implicit measures reflect multiple interacting and qualitatively distinct cognitive processes (e.g., Conrey, Sherman, Gawronski, Hugenberg, & Groom, 2005; Krieglmeyer & Sherman, 2012; Payne, 2001; Payne, Hall, Cameron, & Bishara, 2010). Some of these processes are related to the content of currently accessible associations (e.g., Black–unpleasant association activation), whereas other processes constrain or inhibit the influence of currently accessible mental contents. Both types of processes are affected by current context.

The extent to which implicit bias reflects the contribution of accessible associations versus control-oriented processes suggests a novel implication of the BoC perspective: Would BoC predict that the influence of control processes depends on situational factors? Calanchini, Sherman, Klauer, and Lai (2014) demonstrated that some of the control processes that contribute to responses on the Implicit Association Test (IAT) operate similarly across attitude objects, which suggests that control may be relatively invariant across situations. That said, the control processes that influence responses on implicit measures are clearly not situationally independent. For example, the influence of control is diminished when executive resources are scarce, such as when judgments are made under time pressure (e.g., Conrey et al., 2005; Payne, 2001). Future research should investigate the extent to which accessibility and control processes are influenced by other situational factors. In addition, future work could clarify how accessibility and control contribute to the discrepancy between individual level and aggregate level stability of implicit intergroup biases.

On Reliability

Because the low temporal stability of implicit bias measures is a focal point of the BoC proposal (see PV&L's Puzzle 1), it is relevant to consider methodological and/or psychological factors related to measurement instability. PV&L note that test–retest reliability for measures of intergroup bias, such as the IAT and AMP, fluctuate around $r = .4$ (Gawronski, Morrison, Phills, & Galdi, 2017). These reliabilities are low in comparison to explicit measures of bias. However, it may be more apt to compare the test–retest reliability of implicit measures of bias against measures that are procedurally similar, rather than against procedurally dissimilar explicit measures of bias

(e.g., self-report, feeling thermometers). Examining the reliability of cognitive processing measures, Wöstmann et al. (2013) found the test–retest reliability of the Stroop color-word task to be $r = .70$ and the reliability of the Eriksen Flanker task to be $r = .53$. Measures that rely on rapid sequential priming, such as semantic priming tasks, have shown even lower test–retest reliability. For example Stolz, Besner, and Carr (2005) found strong semantic priming effects at the group level but test–retest reliability of only $r = .4$, even within a single session. Test–retest reliability decreased as measurement conditions became more "automatic" (i.e., shorter vs. longer stimulus-onset asynchrony, smaller vs. larger semantic relatedness proportion). Evidence from this line of work led Stolz and colleagues to conclude that the contents of semantic memory "are inherently noisy and uncoordinated in their automatic or reflexive activities" (p. 328). To emphasize the similarity of this work, 81% of participants showed semantic priming effects under the most constrained measurement conditions, paralleling the 85% of participants who generally show intergroup biases on implicit measures. Yet, simultaneously, test—retest reliability of the semantic priming effect was $r = -.06$. Extrapolating from this, it appears that group level effects can be quite replicable even in the presence of large intraindividual variability, as long as there exists small but reliable biases among most respondents.

Very recent work by Hedge, Powell, and Sumner (2017) describes a paradox similar to that illustrated by PV&L (see Puzzle 1). According to the Reliability Paradox, measures of cognitive processing (e.g., Stroop, Flanker, go/no-go) produce replicable group level estimates but are ineffective tools at capturing differences between participants. Hedge and colleagues found that common measures of cognitive processing could not reliably rank participants' performance. In light of this finding, Hedge and colleagues considered the possibility that individuals' cognitive processing is context specific and highly variable over time—a proposal that maps closely onto that of PV&L's BoC perspective—but reject it in favor of the argument that measures of cognitive processing are inherently limited in their ability to characterize individual differences in processing. In other words, Hedge et al. (2017) maintained that cognitive processing is an individual level construct but conceded that cognitive processing measures appear to produce limited inter-individual variance. The lack of reliability in implicit measures of bias may not reflect the absence of individual level attitudes but instead may reflect the fact that measurement conditions characterizing implicit measures attenuate individual differences in bias. This possibility has serious implications for our interpretations of implicit measures and highlights the necessity of future work to determine the extent to which currently used implicit measures are capable of reliably estimating individual differences in bias.

Summary

PV&L present the BoC model to reconcile the apparent contradiction of stable implicit intergroup biases at the group level and unstable implicit intergroup biases at the individual level. BoC is an intriguing proposal that highlights the importance of

situational influences on implicit bias. However, we have several reservations with the BoC perspective.

First, it is unclear whether BoC is generative above and beyond traditional attitude models that posit that both temporary and chronic accessibility influence the measurement of mental concepts. BoC is unique in its suggestion that implicit measures reflect solely the concepts made temporarily accessible by the perceiver's current situation and not at all by the influence of stable, mental representations developed over the perceiver's learning history. Standard models of chronic and temporary accessibility can reproduce such outcomes by simply altering the relative weights given to each, without entirely dispensing with the notion of (semi-) stable representations.

Second, BoC illustrates one possible means by which formalized intergroup discrimination may be propagated, even while individuals do not endorse discrimination. However, though BoC generates new predictions for the stability of group level intergroup biases, it does not offer an account of how bias got into the system in the first place. The relationship between the personal and the social would seem to be bidirectional, rather than flowing only from the social to the personal.

Third, though BoC emphasizes the role of concept accessibility, context may also influence expressed implicit bias via other mechanisms, including some related to cognitive control. The extent to which each affects contextualized implicit attitudes is an important question for future research.

Finally, care is required when interpreting the reliability of implicit evaluative measures. Reliability is indeed less than ideal from a statistical perspective, and perhaps lower than expected given prominent descriptions of implicit attitudes as stable and resistant to situational influence. Even so, we should consider what benchmarks are appropriate when evaluating the reliability of implicit measures. Measures of cognitive processing also demonstrate surprisingly low test–retest reliability and yet cognitive processing is not questioned as a construct residing within-persons. Similarly, we think it is premature to wholly reject implicit intergroup bias as a construct residing within individuals.

Funding

References

Blair, I. V. (2002). The malleability of automatic stereotypes and prejudice. *Personality and Social Psychology Review*, 6, 242–261. doi:10.1207/S15327957PSPR0603_8

Calanchini, J., Hehman, E., & Salerno, J. (2017). *Regional implicit racial bias predicts wrongful conviction of Black defendants*. Unpublished manuscript.

Calanchini, J., Sherman, J. W., Klauer, K. C., & Lai, C. K. (2014). Attitudinal and non-attitudinal components of IAT performance. *Personality and Social Psychology Bulletin*, 40(10), 1285–1296. doi:10.1177/0146167214540723

Calanchini, J., Witkowski, P., Sparks, J., Schmidt, K., & Sherman, J. W. (2017). *County level implicit gender stereotypes predict opposition to*

Hillary Clinton in the 2016 American presidential election*. Unpublished manuscript.

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

Conrey, F. R., Sherman, J. W., Gawronski, B., Hugenberg, K., & Groom, C. J. (2005). Separating multiple processes in implicit social cognition: The quad model of implicit task performance. *Journal of Personality and Social Psychology*, 89(4), 469–487. doi:10.1037/0022-3514.89.4.469

Conrey, F. R., & Smith, E. R. (2007). Attitude representation: Attitudes as patterns in a distributed, connectionist representational system. *Social Cognition*, 25, 739–758. doi:10.1521/soco.2007.25.5.718

Cronbach, L. J. (1957). The two disciplines of scientific psychology. *American Psychologist*, 12, 671–684. doi:10.1037/h0043943

Devine, P. G. (1989). Stereotypes and prejudice: Their automatic and controlled components. *Journal of Personality and Social Psychology*, 56, 5–18. doi:10.1037/0022-3514.56.1.5

Dunham, Y., Chen, E. E., & Banaji, M. R. (2013). Two signatures of implicit intergroup attitudes: Developmental invariance and early enculturation. *Psychological Science*, 26(6), 860–868. doi:10.1177/0956797612463081

Fazio, R. H. (2007). Attitudes as object-evaluation associations of varying strength. *Social Cognition*, 25, 603–637. doi:10.1521/soco.2007.25.5.603

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(6), 1013–1027. doi:10.1037/0022-3514.69.6.1013

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., 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*.

Greenwald, A. G., Banaji, M. R., Rudman, L. A., Farnham, S. D., Nosek, B. A., & Mellott, D. S. (2002). A unified theory of implicit attitudes, stereotypes, self-esteem, and self-concept. *Psychological Review*, 109, 3–25. doi:10.1037//0033-295X.109.1.3

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

Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of implicit attitudes. *Journal of Experimental Psychology General*, 143, 1369–1392. doi:10.1037/a0035028

Hedge, C., Powell, G., & Sumner, P. (2017). The reliability paradox: Why robust cognitive tasks do not produce reliable individual differences. *Behavior Research Methods*. Advance online publication. doi:10.3758/s13428-017-0935-1

Hehman, E., Calanchini, J., Flake, J., & Leitner, J. (2017). *The construct validity of region level implicit and explicit racial bias*. Unpublished manuscript.

Hehman, E., Flake, J.K., & Calanchini, J. (in press). Disproportionate use of lethal force in policing is associated with regional racial biases of residents. *Social Psychological and Personality Science*. Advance online publication. doi:10.1177/1948550617711229

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.

Krieglmeyer, R., & Sherman, J. W. (2012). Disentangling stereotype activation and stereotype application in the stereotype misperception task. *Journal of Personality and Social Psychology*, 103(2), 205–224. doi:10.1037/a0028764

Lai, C. K., Skinner, A. L., Cooley, E., Murray, 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. doi:10.1037/xge0000179

Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016a). Blacks death rate due to circulatory diseases is positively related to Whites explicit racial bias: A nationwide investigation using Project Implicit. *Psychological Science*, *27*(10), 1299–1311. doi:10.1177/0956797616658450

Leitner, J.B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016b). 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

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

Obschonka, M., Stuetzer, M., Audretsch, D. B., Rentfrow, P. J., Potter, J., & Gosling, S. D. (2016). Macropsychological factors predict regional economic resilience during a major economic crisis. *Social Psychological and Personality Science*, *7*(2), 95–104. doi:10.1177/1948550615608402

Olson, M. A., & Fazio, R. H. (2004). Reducing the influence of extrapersonal associations on the Implicit Association Test: Personalizing the IAT. *Journal of Personality and Social Psychology*, *86*, 653–667. doi:10.1037/0022-3514.86.5.653

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. (2001). Prejudice and perception: The role of automatic and controlled processes in misperceiving a weapon. *Journal of Personality and Social Psychology*, *81*, 181–192. doi:10.1037/0022-3514.81.2.181

Payne, B. K., Hall, D., Cameron, C. D., & Bishara, A. J. (2010). A process model of affect misattribution. *Personality and Social Psychological Bulletin*, *36*, 1397–1408. doi:10.1177/0146167210383440

Ranganath, K. A., Smith, C. T., & Nosek, B. A. (2008). Distinguishing automatic and controlled components of attitudes from direct and indirect measurement methods. *Journal of Experimental Social Psychology*, *44*(2), 386–396. doi:10.1016/j.jesp.2006.12.008

Rentfrow, P. J., Gosling, S. D., Jokela, M., Stillwell, D. J., Kosinski, M., & Potter, J. (2013). Divided we stand: Three psychological regions of the United States and their political, economic, social, and health correlates. *Journal of Personality and Social Psychology*, *105*(6), 996–1012. doi:10.1037/a0034434

Stolz, J. A., Besner, D., & Carr, T. H. (2005). Implications of measures of reliability for theories of priming: Activity in semantic memory is inherently noisy and uncoordinated. *Visual Cognition*, *12*(2), 284–336. doi:10.1080/13506280444000030

Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, *107*, 101–126. doi:10.1037/0033-295X.107.1.101

Wöstmann, N. M., Aichert, D. S., Costa, A., Rubia, K., Möller, H. J., & Ettinger, U. (2013). Reliability and plasticity of response inhibition and interference control. *Brain and Cognition*, *81*(1), 82–94. doi:10.1016/j.bandc.2012.09.010

Zerhouni, O., Rougier, M., & Muller, D. (2016). "Who (really) is Charlie?": French cities with lower implicit prejudice toward Arabs demonstrated larger participation rates in Charlie Hebdo rallies. *International Review of Social Psychology*, *29*, 69–76. doi:10.5334/irsp.50