Association for Psychological Science

Perspectives on Psychological Science
2022, Vol. 17(2) 611–613

DOI: 10.1177/17456916211070833
www.psychologicalscience.org/PPS
SAGE

A Strange Kind of Wave: Response to Payne, Vuletich, and Lundberg (2022)

Paul Connor
Department of Psychology, New York University

I am pleased to have the chance to respond to Payne et al.'s (2022) commentary on Connor and Evers (2020), which was itself a critique of their 2017 bias-of-crowds theory article about the nature of implicit bias (Payne et al., 2017). To their credit, Payne and colleagues' commentary has clarified a number of issues, so I will make this response brief.

Clarification 1: The Puzzles

A first clarification provided by Payne and colleagues' (2022) commentary regards the empirical puzzles they described with regard to implicit bias. Each of these puzzles referred to the observation of higher correlations (test–retest correlations, correlations with relevant criteria, etc.) at aggregate levels (e.g., state, county, country) compared with the individual level. They wrote:

> For most psychological measures, we would expect stable means because of stable individual scores. But the low stability estimates imply that the rank orders of participants (i.e., who is high on implicit prejudice and who is low) change dramatically from one measurement occasion to another, whereas the group mean scores somehow remain constant. Replicable implicit bias on average, paired with constantly shifting individual scores, presents an important puzzle to be explained. (Payne et al., 2017, p. 234)

In our critique, we used simulated data to demonstrate that these observations are in fact relatively routine results of measurement error, nonzero group-level variation, and aggregation. We concluded:

> Given levels of measurement accuracy and [intraclass correlation coefficients] similar to those observed in the case of implicit bias, aggregation of noisy individual-level scores produced both greater stability of group means compared with individual scores . . . and greater correlations with related criteria . . . at the group level than at the individual level. Payne and colleagues' empirical puzzles are therefore not puzzling at all. (Connor & Evers, 2020, p. 9)

The authors' commentary clarifies that they agree with us on this. Discussing our simulated data, they write:

> This simulation provides evidence consistent with the view that implicit bias is a noisily measured construct at the individual level, which becomes less noisily measured and more strongly correlated with criterion variables when measured in the aggregate. Because this "alternative view" is identical to that posited by the bias-of-crowds model, we have nothing to dispute here. (Payne et al., 2022, p. 607)

Everyone therefore agrees, it seems, that the puzzles are puzzling only insofar as one fails to appreciate the expected effects of aggregation in the presence of measurement error and nonzero systematic group-level variation. I consider it a legitimate question whether it was ever appropriate to describe such phenomena as "important problems for the science of implicit bias to solve" (Payne et al., 2017, p. 235), given that they are both easily explained and not specific to implicit bias, but that is a subjective matter.

Clarification 2: Most Systematic Variance in Implicit Bias

A second clarification concerns Payne and colleagues' claim that high aggregate-level correlations substantiate the claim that "most of the systematic variance in implicit bias is situational" (Payne et al., 2017, p. 234).

Corresponding Author:
Paul Connor, Department of Psychology, New York University
Email: pconnor@berkeley.edu

In our critique, we disagreed with this, and pointed out that large aggregate-level correlations can in fact represent negligible amounts of overall variation in a measure. For example, aggregating implicit association test (IAT) scores within weekdays produces a near perfect test–retest correlation (r = .95), despite the fact that weekdays explain just 0.01% of overall variance in IAT scores (Connor & Evers, 2020). By contrast, despite relatively lower test–retest correlations at the individual level (a recent meta-analytic estimate of the individual-level test–retest reliability of the IAT was .50; Greenwald & Lai, 2020), approximately 50% of overall variance in IAT scores can be accounted for by stable individual differences—around 5,000 times more than weekdays.[1] In their commentary, Payne and colleagues appear to have taken this criticism on board and have offered a revised wording of their original claim:

> The issue of variability at different levels of analysis highlights a statement in our original article that we now see was ambiguous and may be a cause of confusion. We wrote that "most of the systematic variance in implicit biases appears to operate at the level of situations" (p. 236). Our intended claim was not that there is *more variance* between situations than between individuals. Our intended claim was that the variance between situations is *more systematic* than the variance between individuals. By "more systematic," we mean more reliable and valid, as evidenced by test–retest reliability and correlations with criterion variables. (Payne et al., 2022, p. 607)

We therefore agree that aggregate-level measurement can achieve higher test–retest and criterion correlations than individual-level measurement, despite that fact that there will generally be greater overall amounts of systematic variance overall at the individual level. And although I worry that readers will continue to be confused by the distinction between "more total systematic variance" and "variance that is more systematic,"' I hope the present exchange has helped clear this up to some degree.

Remaining Disagreements

Some remaining disagreements between Payne and colleagues' views and my own deserve mention. First, I am skeptical of Payne and colleagues' claim that their model provides a valuable new perspective that "raises new and different research questions" (Payne et al., 2022, p. 606). As far as I can see, researchers have never doubted that implicit bias would vary systematically across contexts and began studying aggregate-level questions about implicit bias as soon as doing so became possible, as a result of the existence and growing awareness of the Project Implicit database (Xu et al., 2014). For example, before the model was introduced, my coauthors and I used IAT scores aggregated at the state-year level as a measure of regional racial bias (this project began in 2015 and was eventually published in Connor et al., 2019). Doing so, we were inspired by other pre–bias-of-crowds work by colleagues who aggregated IAT scores at the state level (Leitner et al., 2016). I therefore do not think that the bias-of-crowds model was or is necessary to inspire aggregate-level questions about implicit bias.

A second and perhaps more interesting disagreement between Payne and colleagues and me concerns the plausibility of their broader overall vision of what implicit bias is. To recap, the authors argued that implicit bias should be thought of as being primarily a feature of social environments, reified within individuals from moment to moment because of exposure to contextual structural inequalities. They likened this process to how "the wave" passes through sports fans at a stadium (Payne et al., 2017).

The most obvious problem with this metaphor has been covered: 50% of variance in IAT scores is accounted for by stable individual differences. This is therefore a very strange kind of "wave," in which we can in theory predict with reasonable success whether individuals will be sitting or standing at any time from their previous behavior. Personally, I find it difficult to visualize such a wave; the necessary mental image involves thousands of sports fans entering and exiting a stadium frozen semipermanently in crouched or seated positions.

A more interesting problem with this vision, however, is simply that it does not fit particularly well with everything we know about the phenomenon of implicit bias. Payne and colleagues described how they believe structural inequalities create and maintain implicit biases:

> We assume that individuals have a variety of accessible links to social categories, many of which are fleeting and changeable from one context to another, which makes individual scores variable and unreliable across time. The robust average effects arise because there is a high degree of prejudice in contemporary culture, communicated in countless ways, from biased depictions in media, to segregation in everyday interactions, to daily observations of which social groups tend to occupy high-status and low-status positions, and so on. Inequality in the culture exerts a constant influence that gives rise to large average effects, like a slow moving wave across generations. (Payne et al., 2017, p. 238)

Admittedly, it seems compelling to link implicit racial bias—the primary focus of Payne and colleagues' work on the bias of crowds—to structural features of the environment. And this is especially the case in the United States, given its pervasive and well-documented racial inequality. But it is also important to remember that implicit biases are not limited only to race. Individuals display implicit evaluative preferences with regard to a wide variety of conceptual categories, including implicit preferences for female over male, for the future over the past, for summer over winter, for vegetables over meat, for Coke over Pepsi, and for feminism over traditional values (Nosek, 2005).

So, I am curious: Do the authors of the bias-of-crowds model believe that all of these biases result from exposure to structural inequalities? How do we make sense, for example, of the observed implicit preference for feminism over traditional values? Does inequality in the culture also reify this bias from moment to moment in individuals' minds? Perhaps it might be argued that some implicit biases are not perpetually reified by structural inequalities in the environment and simply reflect individuals' learned preferences and semantic associations, which makes it slightly easier, for example, to perform categorization tasks when the same computer key is assigned for both "feminism" and "good" or for both "meat" and "bad." But if this is possible in the case of some biases, I do not see why this could this not be the case for all implicit biases. Payne and colleagues would therefore benefit, I believe, from broadening their focus and considering how their model might account for other kinds of implicit bias that do not correspond in obvious ways with structural inequalities.

Conclusion

Ultimately, I expect that these kinds of problems will likely do little to hamper enthusiasm for the bias-of-crowds model, which in just 4 years has amassed hundreds of citations. However, beyond a certain amount of confusion and misapprehension about the nature of implicit bias, I also expect that little harm will come of this popularity. At its core, the primary claim of the model is an uncontroversial one: There is systematic group-level variation in implicit bias. Everything else, such as the empirical puzzles and their solutions, is a natural consequence of this. To my knowledge, no one has ever doubted the existence of some level of systematic group-level variation in implicit bias. And researchers—myself included—have been aggregating implicit bias data at the group level and thinking about aggregate-level questions since well before the bias-of-crowds model was conceived. I believe that when done well, such work can contribute positively to human knowledge, and that is a goal I wholeheartedly share and support.

Transparency

Action Editor: Laura A. King
Editor: Laura A. King
Declaration of Conflicting Interests
The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

ORCID iD

Paul Connor https://orcid.org/0000-0002-4995-4679

Note

1. An average test–retest correlation of $r = 0.50$ implies a correlation of $r = 0.71$ between individuals' true, stable levels of bias and measured IAT scores.

References

Connor, P., & Evers, E. R. (2020). The bias of individuals (in crowds): Why implicit bias is probably a noisily measured individual-level construct. *Perspectives on Psychological Science*, *15*(6), 1329–1345. https://doi.org/10.1177/1745691620931492

Connor, P., Sarafidis, V., Zyphur, M. J., Keltner, D., & Chen, S. (2019). Income inequality and White-on-Black racial bias in the United States: Evidence from project implicit and Google trends. *Psychological Science*, *30*(2), 205–222. https://doi.org/10.1177/0956797618815441

Greenwald, A. G., & Lai, C. K. (2020). Implicit social cognition. *Annual Review of Psychology*, *71*, 419–445. https://doi.org/10.1146/annurev-psych-010419-050837

Leitner, J. B., Hehman, E., Ayduk, O., & Mendoza-Denton, R. (2016). 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. https://doi.org/10.1177/0956797616658450

Nosek, B. A. (2005). Moderators of the relationship between implicit and explicit evaluation. *Journal of Experimental Psychology: General*, *134*(4), 565–584. https://doi.org/10.1037/0096-3445.134.4.565

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*(4), 233–248. https://doi.org/10.1080/1047840X.2017.1335568

Payne, B. K., Vuletich, H. A., & Lundberg, K. B. (2022). Critique of the bias-of-crowds model simply restates the model: Reply to Connor and Evers (2020). *Perspectives on Psychological Science*, *17*(2), 606–610. https://doi.org/10.1177/1745691621997394

Xu, K., Nosek, B., & Greenwald, A. (2014). Psychology data from the race implicit association test on the project implicit demo website. *Journal of Open Psychology Data*, *2*(1), Article e3. https://doi.org/10.5334/jopd.ac