

# I Saw You in the Crowd: Credibility, Reproducibility, and Meta-Utility

Nate Breznau, *University of Bremen*

Crowdsourcing is a new word. Google Books' first recorded usage of "crowdsourcing" was in 1999.<sup>1</sup> It originally referenced Internet users adding content to websites, the largest example being Wikipedia (Zhao and Zhu 2014). As a scientific practice, its roots go back to open-engineering competitions in the 1700s, when the collective ideas of many overcame otherwise insurmountable problems of the few. Openly tapping into a large population breaks down barriers of time, money, and data. The solution is simple. A centralized investigator or team poses a question or problem to the crowd. Sometimes tasks are co-creative, such as the programming of packages by statistical software users and the development of Wiki content. Other times, they are discrete, such as asking individuals to donate computing power,<sup>2</sup> respond to a survey, or perform specific tasks. Using crowdsourced methods, scientists, citizens, entrepreneurs, and even governments more effectively address societies' most pressing problems (e.g., cancer and global warming) (Howe 2008).

Although the public comprises the typical crowd, new landmark studies involved researchers crowdsourcing other researchers. The human genome project, for example, involved major byline contributions from almost 3,000 researchers. In the social sciences, crowdsourcing of researchers is brand new. In a 2013 study, Silberzahn et al. (2018) convened researchers from across disciplines and geographic locations to analyze the same dataset to discover whether football referees issued more red cards to darker-skinned players.<sup>3</sup> Klein et al. (2014) sourced laboratories across the world to try and reproduce several high-profile experimental psychological studies. In the *Crowdsourced Replication Initiative*, Breznau, Rinke, and Wuttke et al. (2019) demonstrated that crowdsourcing also is useful for replication, structured online deliberation, and macro-comparative research. Studies like these are a new paradigm within social, cognitive, and behavioral research.

In the knowledge business, crowdsourcing is highly cost effective. The average researcher lacks the means to survey or conduct experiments with samples outside of universities. Crowdsourcing platforms such as Amazon's Mechanical Turk (mTurk) changed this by creating access to even-more-convenient convenience samples from all over the world (Palmer and Strickland 2016). Others use these platforms to crowdsource globally competitive labor to perform paid tasks

of scientific value (Berinsky, Huber, and Lenz 2012). Blockchain technology uses decentralized crowd computing. Many actions that crowds already take (e.g., "Tweeting" and "geotagging" images) offer free possibilities to test hypotheses of social and political relevance (Nagler and Tucker 2015; Salesses, Schechtner, and Hidalgo 2013). Wikipedia had fantastic success, given the strong willingness of publics around the world to add and monitor content. The Wiki model is the basis for all types of crowd creations.<sup>4</sup> There are the Wikiversity<sup>5</sup> crowdsourcing-teaching resources and the Replication-Wiki<sup>6</sup> for listing and identifying replications across the social sciences (Höfler 2017). In a similar vein, there are efforts to crowdsource open peer review and ideas for what researchers should study from potential stakeholders outside of science.

Some scholars might question why many researchers are necessary when we have so many preexisting datasets and a single researcher can run every possible statistical model configuration. Crowdsourcing researchers' time and effort may seem inefficient when machine learning is so advanced. However, if a single scholar or machine throws all possible variables into their models like a "kitchen sink," it is only efficient at maximizing prediction. There is a subtle but crucial difference between human research, where if  $X$  predicts  $Y$ , it *might* be a meaningful association, versus the machine approach, where if  $X$  predicts  $Y$ , it *is* meaningful (Leamer 1978). When humans or machines start running every possible model, which is the typical machine-learning strategy, they implicitly test every possible hypothesis and potentially every possible theory. Among all possible models, there will be some in which the data-generating process is theoretically impossible, such as predicting biological sex as an outcome of party affiliation, or right-party votes in 1970 from GDP in 2010. If we want something more than predictive power, human supervision and causal inference are necessary to rule out impossible realities that a computer cannot identify on its own (Pearl 2018).

Crowdsourcing is epistemologically different than machine learning. It is not a "kitchen-sink" method. The landmark Silberzahn et al. (2018) study demonstrated that seemingly innocuous decisions in the data-analysis phase of research can change results dramatically. The effect of idiosyncratic features of researchers (e.g., prior information and beliefs) are something Bayesian researchers raised warning signals about long ago (Jeffreys 1998 [1939]). This is echoed in more recent

discussions about the unreliability and irreproducibility of research (Dewald, Thursby, and Anderson 1986; Gelman and Loken 2014; Wicherts et al. 2016). However, grappling with this uncertainty is problematic. Identification of informative research priors usually involves consulting an expert, such as the researcher doing the work, or using simulations wherein the user selects parameters that may (or may not) reflect the

remains to be seen how valuable the machines are with minimal human interference.

At the heart of the SCORE project are questions of credibility and reliability. These topics also are at the center of scandals and conflicts that are infecting science at the moment and rapidly increasing researchers and funding agencies' interests in replication (Eubank 2016; Ishiyama 2014; Laitin

### *Crowdsourcing is epistemologically different than machine learning. It is not a "kitchen-sink" method.*

reality of the researcher's decisions. The primary objection is that the distribution of prior beliefs across researchers is unknown (Gelman 2008). The Silberzahn et al. (2018) and Breznau, Rinke, and Wuttke (2018; 2019) studies provide a *real* observation of these priors because all researchers develop what should be at least a plausible model for testing the hypothesis, given the data. This observation extends beyond prior studies because the crowdsourced researchers scrutinized the other researchers' models before seeing the results. This practice is extremely useful because it is not one expert guessing about the relative value of a model but rather potentially 100 or more.

The reduction of model uncertainty is good from a Bayesian perspective but it also demands theory from a causal perspective. Statisticians are well aware that without the correct model specification, estimates of uncertainty are themselves uncertain, if not useless. Thus, rather than having 8.8 million false positives after running 9 billion different regression models (Muñoz and Young 2018), humans can identify correct—or at least “better”—model specifications using causal theory and logic (Clark and Golder 2015; Pearl 2018). The diversity of results in the Silberzahn et al. (2018) and Breznau, Rinke, and Wuttke et al. (2018; 2019) studies helped to identify key variables and modeling strategies that had the “power” to change the conclusions of any given research team. These were not simply variables among millions of models but rather variables among carefully constructed plausible models. This shifts the discussion from the results and even the Bayesian priors to data-generating theories.

Machines cannot apply discriminating logic—or can they? The US military's Defense Advanced Research Projects Agency (DARPA) deems the human-versus-machine question so important that it currently funds the Systematizing Confidence in Open Research and Evidence (SCORE) project pitting machines against research teams in reviewing the credibility of research across disciplines. SCORE architects believe that machines might be capable of determining the credibility of social and behavioral research better, faster, and more cost effectively than humans. In this case, crowdsourcing provides DARPA the exact method it needs to answer this question. This project is the largest crowdsourcing of social researchers to date, with already more than 500 participating.<sup>7</sup> It is a major demonstration that crowdsourcing can bring together the interests of academics and policy makers. It

and Reich 2017; Stockemer, Koehler, and Lentz 2018). However, there are limits to what replications can offer, and they take on diverse formats and goals (Clemens 2015; Freese and Peterson 2017). Currently, the decision of whether, what, and how to replicate is entirely a researcher's prerogative. Few engage in any replications, meaning that any given original study is fortunate to have even one replication.<sup>8</sup> Journals thus far have been hesitant to publish replications, especially of their own publications, even if the replications overturn preposterous-sounding claims (e.g., precognition) (Ritchie, Wiseman, and French 2012) or identify major mistakes in the methods (Breznau 2015; Gelman 2013). Crowdsourcing could change this because it provides the power of meta-replication in one study. Moreover, crowdsourcing might appear to journal editors as cutting-edge research rather than “just another” replication (Silberzahn and Uhlmann 2015).

Perhaps more important, crowdsourced replications provide reliability at a meta-level. Replications, experimental reproductions, and original research alike suffer from what Leamer (1978) referred to as “metastatistics” problems. If researchers exercise their degrees of freedom such that ostensibly identical research projects have different results—for example, more than 5% of the time—then any one study is not reliable by most standards. Breznau, Rinke, and Wuttke (2018) simulated this problem and deduced that four to seven independent replications are necessary to obtain a majority of direct replications—that is, a reproduction of the same results, arriving at similar effect sizes within 0.01 of an original study in a 95% confidence interval. Subsequently, the results from their crowdsourced project showed that even after correcting for major mistakes in code, 11% of the effects were substantively different from the original (i.e., a change in significance or direction) (Breznau, Rinke, and Wuttke 2019). Søndergaard (1994) reviewed replications of the Values Survey Model employed by Hofstede—one of the most cited and replicated studies in social science—and found that 19 of 28 replications (only 68%) came to approximately the same relative results. Lack of reproducibility is not restricted to quantitative research. For example, Forscher et al. (2019) investigated thousands of National Institutes of Health grant reviews and found that using the typical three to five reviewers approach achieved only a 0.2 inter-rater reliability. Increasing this to an extreme of 12 reviewers per application achieved only 0.5 reliability. Both scores are far below any acceptable standard for researchers subjectively evaluating the same text data.

There is much to learn about why researchers arrive at different results. It appears that intentional decisions are only one factor (Yong 2012) and that context, versioning, and institutional constraints also play a role (Breznau 2016; Breznau, Rinke, and Wuttke 2019). Crowdsourcing addresses this meta-uncertainty because principal investigators (PIs) can hold certain research factors constant—including the method, data, and exact framing of the hypothesis—thereby exponentially increasing the power to both discover why researchers' results differ and to meta-analyze a selected topic (Uhlmann et al. 2018). Combined with the rapidly expanding area of specification-curve analysis, crowdsourcing provides a new way to increase credibility for political and social research—in both sample populations and among the researchers themselves (Rohrer, Egloff, and Schmukle 2017; Simonsohn, Simmons, and Nelson 2015). It is hoped that these developments are tangible outcomes that increase public, private, and government views of social science.

As social scientists, we naturally should be skeptical of the “DARPA hypothesis” that computers could be more reliable than humans for evidence-based policy making (Grimmer 2015). In fact, a crowdsourced collaboration of 107 teams recently demonstrated that humans were essentially as good as machines in predicting life outcomes of children (Salganik et al. 2020). In this study, however, both humans and machines were relatively poor at prediction overall. These problems may stem from a lack of theory. For example, in a different study of brain imaging involving 70 crowdsourced teams, not a single pair agreed about the data-generating model (Botvinik-Nezer et al. 2020). Breznau et al. (2019) suggested that a lack of theory to describe the relationship of immigration and social policy preferences means that research in this area also is unreliable and—as these new crowdsourced studies demonstrate—meta-analyzing the results does not solve this dearth of theory. However, crowdsourcing has an untapped resource in addition to meta-analysis of results for resolving these issues: meta-construction of theory (Breznau 2020). In the Breznau et al. (2019) study, careful consideration of the data-generating model, online deliberation, and voting on others' models revealed where theoretical weaknesses and lack of consensus exist. When coupling observed data-generating model variations across teams with their deliberations and disagreements on the one hand, with how these may or may not shape the results on the other, immigration and social policy scholars gain insights to where they should focus their theoretical efforts in the future to obtain the most significant gains in knowledge.

Additionally, as a structured crowdsourced research endeavor using the technology of the online deliberation platform Kialo, the Breznau et al. (2019) study undermined the current research system that favors novelty of individual researchers. It instead promoted consensus building and direct responsiveness to theoretical claims. When else do hundreds of political researchers collaborate to focus their theoretical discussions in one area? Crowdsourcing, with the help of technologies such as Kialo, could resolve the perpetual problem of scholars, areas, and disciplines talking “past” one another. It also would provide a leap forward technologically:

we currently observe “collective” theory construction in a primitive format among conference panels and journal symposia, limited to a few invited participants in a specific event in space and time. This is a narrow collaborative process in the face of structured crowdsourced deliberations in which potentially thousands of global participants can engage at their convenience over many months. Kialo also provides extensive data for analyzing deliberative crowd processes and outcomes, thereby streamlining the process.<sup>9</sup>

The benefits of crowdsourcing are not only on the scientific output or theoretical side. The Silberzahn et al. (2018) and Breznau et al. (2018; 2019) studies also included structured interaction among researchers, which fostered community engagement and the exchange of ideas and methodological knowledge. The Breznau et al. (2019) study specifically asked participants about their experiences and found that learning and enjoyment were common among them. For example, 69% reported that participation was enjoyable, 22% were neutral, and only 7% found it not enjoyable. Of 188 participants, the retention rate was 87% from start to finish, which demonstrates motivation across the spectrum of PhD students, postdocs, professors, and nonacademic professionals. Crowdsourcing potentially requires a significant time investment on behalf of the participating researchers, but the investment pays off in academic and personal development. Moreover, the practice of crowdsourcing researchers embodies many principles amenable to the Open Science Movement, such as open participation, transparency, reproducibility, and better practices. Crowdsourcing may not be a high-density method but it seems sustainable as an occasional large-scale project and as something good for science.

A key principle of open science and crowdsourcing also is related to incentive structure. The status attainment of scholars is theoretically infinite, limited only by the number of people willing to cite a scholar's work, which leads to egomaniacal and destructive behaviors (Sørensen 1996). This is the status quo of academia. In a crowdsourced endeavor, the crowdsourced researcher participants are equal. There is no chance for cartelism, veto, or exclusionary practices as long as the PIs are careful moderators. Thus, in its ideal form, crowdsourced projects should give authorship to all participants who complete the tasks assigned to them. The PIs of a crowdsourced endeavor naturally have more to gain than the participants because the research is their “brainchild” and they will be known as the “PIs.” Nonetheless, they also have more to lose in case such a massive implementation of researcher human capital were to fail; therefore, the participants have good reason to trust the process. The net gains should be positive for everyone involved in a well-executed crowdsourced project. The advantage of this vis-à-vis the Open Science Movement is that the current “publish-or-perish” model becomes something positive, wherein all participants benefit by getting published—and any self-citation thereafter is a community rather than individualistic good.

Crowdsourcing is not without boundaries. There is a limited supply of labor. Crowdsourcing of a replication or original research essentially asks researchers to engage in a project that may require as much empirical work as one of their

“own” projects. Moreover, the newness of this method means that researchers may not see the value in investing their time. Before contributing their valuable time, researchers may want to consider (1) the novelty and importance of the topic, and (2) whether the proposed project fits with open and ethical science ideals. In the case of crowdsourced replications, the “topic” is to end the replication crisis. Bates (2016) suggested a formula for selecting studies to replicate. He argued that priority for ending the replication crisis can be measured as influence (i.e., citations) divided by evidence (i.e., number of replications, wherein zero might logically be 0.01). The larger the number, the higher the replication priority. I suggest extending this formula to include sociopolitical relevance of the topic, something that crowdsourcing can be used to measure. Crowdsourced ranking of studies in terms of impact can be a weighting coefficient on the right-hand side of this equation. This would provide the possibility to end the replication crisis and also solve pressing societal problems. ■

#### NOTES

1. As of August 6, 2020.
2. For example, see [www.worldcommunitygrid.org](http://www.worldcommunitygrid.org) or cryptocurrency using blockchain technology.
3. Also known as “soccer” in a minority of countries.
4. Just check the growing list on the Wikipedia crowdsourcing projects page! Available at [https://en.wikipedia.org/wiki/List\\_of\\_crowdsourcing\\_projects](https://en.wikipedia.org/wiki/List_of_crowdsourcing_projects).
5. See [https://en.wikiversity.org/wiki/Teaching\\_and\\_Learning\\_Online](https://en.wikiversity.org/wiki/Teaching_and_Learning_Online).
6. See [http://replication.uni-goettingen.de/wiki/index.php/Main\\_Page](http://replication.uni-goettingen.de/wiki/index.php/Main_Page).
7. See [www.darpa.mil/program/systematizing-confidence-in-open-research-and-evidence](http://www.darpa.mil/program/systematizing-confidence-in-open-research-and-evidence).
8. Hofstede’s Values Survey, the observer effect, and the backfire effect are notable exceptions.
9. This does not mean that Kialo is the only option; however, it seems to be one of the best suited for crowdsourcing.

#### REFERENCES

- Bates, Timothy. 2016. “What Should We Replicate? Things That Will End the Replication Crisis.” *Medium* op-ed. Available at <https://medium.com/@timothycbates/what-should-we-replicate-things-that-will-end-the-replication-crisis-cb09ce24b25f>.
- Berinsky, Adam J., Gregory A. Huber, and Gabriel S. Lenz. 2012. “Evaluating Online Labor Markets for Experimental Research: Amazon.Com’s Mechanical Turk.” *Political Analysis* 20 (3): 351–68.
- Botvinik-Nezer, Rotem, Felix Holzmeister, Colin F. Camerer, Anna Dreber, Juergen Huber, Magnus Johannesson, Michael Kircher, et al. 2020. “Variability in the Analysis of a Single Neuroimaging Dataset by Many Teams.” *Nature* 582 (7810): 84–88.
- Breznau, Nate. 2015. “The Missing Main Effect of Welfare-State Regimes: A Replication of ‘Social Policy Responsiveness in Developed Democracies’ by Brooks and Manza.” *Sociological Science* 2:420–41.
- Breznau, Nate. 2016. “Secondary Observer Effects: Idiosyncratic Errors in Small-N Secondary Data Analysis.” *International Journal of Social Research Methodology* 19 (3): 301–18.
- Breznau, Nate. 2020. “Meta-Construction of Social Theory.” *Crowdid*. Available at <https://crowdid.hypotheses.org/372>.
- Breznau, Nate, Eike Mark Rinke, and Alexander Wuttke. 2018. “Pre-Registered Analysis Plan for ‘How Many Replicators’ Experiment.” Mannheim Centre for European Social Research. *SocArXiv*. Available at <https://osf.io/hkpd>.
- Breznau, Nate, Eike Mark Rinke, Alexander Wuttke, Muna Adem, Jule Adriaans, Amalia-Alvarez-Benjumea, Henrik Andersen, et al. 2019. *Crowdsourced Replication Initiative: Executive Report*. Mannheim Centre for European Social Research. *SocArXiv*. Available at <https://osf.io/preprints/socarxiv/6j9qb>.
- Clark, William Roberts, and Matt Golder. 2015. “Big Data, Causal Inference, and Formal Theory: Contradictory Trends in Political Science? Introduction.” *PS: Political Science & Politics* 48 (1): 65–70.
- Clemens, Michael A. 2015. “The Meaning of Failed Replications: A Review and Proposal.” *Journal of Economic Surveys* 31 (1): 326–42.
- Dewald, William G., Jerry G. Thursby, and Richard G. Anderson. 1986. “Replication in Empirical Economics: The Journal of Money, Credit and Banking Project.” *American Economic Review* 76 (4): 587–603.
- Eubank, Nicholas. 2016. “Lessons from a Decade of Replications at the *Quarterly Journal of Political Science*.” *PS: Political Science & Politics* 49 (2): 273–76.
- Forscher, Patrick, Markus Brauer, William Cox, and Patricia Devine. 2019. “How Many Reviewers Are Required to Obtain Reliable Evaluations of NIH R01 Grant Proposals?” *PsyArXiv* preprint. Available at <https://psyarxiv.com/483zj>.
- Freese, Jeremy, and David Peterson. 2017. “Replication in Social Science.” *Annual Review of Sociology* 43 (1): 147–65.
- Gelman, Andrew. 2008. “Objections to Bayesian Statistics.” *Bayesian Analysis* 3 (3): 445–49.
- Gelman, Andrew. 2013. “Ethics and Statistics: It’s Too Hard to Publish Criticisms and Obtain Data for Republication.” *Chance* 26 (3): 49–52.
- Gelman, Andrew, and Eric Loken. 2014. “The Statistical Crisis in Science.” *American Scientist* 102 (6): 460.
- Grimmer, Justin. 2015. “We Are All Social Scientists Now: How Big Data, Machine Learning, and Causal Inference Work Together.” *PS: Political Science & Politics* 48 (1): 80–83.
- Höfler, Jan. 2017. “Replication Wiki: Improving Transparency in the Social Sciences.” *D-Lib Magazine* 23 (3/4).
- Howe, Jeff. 2008. *Crowdsourcing: How the Power of the Crowd Is Driving the Future of Business*. New York: Random House Business Books.
- Ishiyama, John. 2014. “Replication, Research Transparency, and Journal Publications: Individualism, Community Models, and the Future of Replication Studies.” *PS: Political Science & Politics* 47 (1): 78–83.
- Jeffreys, Harold. 1998 [1939]. *Theory of Probability*. Third edition. New York and Oxford: Clarendon Press and Oxford University Press.
- Klein, Richard A., Kate A. Ratliff, Michelangelo Vianello, Reginald B. Adams, Jr., Štěpán Bahník, Michael J. Bernstein, Konrad Bocian, et al. 2014. “Investigating Variation in Replicability: A ‘Many Labs’ Replication Project.” *Social Psychology* 45 (3): 142–52.
- Latín, David D., and Rob Reich. 2017. “Trust, Transparency, and Replication in Political Science.” *PS: Political Science & Politics* 50 (1): 172–75.
- Leamer, Edward E. 1978. *Specification Searches: Ad Hoc Inference with Nonexperimental Data*. New York: John Wiley & Sons, Inc.
- Muñoz, John, and Cristobal Young. 2018. “We Ran 9 Billion Regressions: Eliminating False Positives through Computational Model Robustness.” *Sociological Methodology* 48 (1): 1–33.
- Nagler, Jonathan, and Joshua A. Tucker. 2015. “Drawing Inferences and Testing Theories with Big Data.” *PS: Political Science & Politics* 48 (1): 84–88.
- Palmer, Joshua C., and Justin Strickland. 2016. “A Beginner’s Guide to Crowdsourcing: Strengths, Limitations, and Best Practices.” *Psychological Science Agenda*. Available at [www.apa.org/science/about/psa/2016/06/changing-minds](http://www.apa.org/science/about/psa/2016/06/changing-minds).
- Pearl, Judea. 2018. “Theoretical Impediments to Machine Learning with Seven Sparks from the Causal Revolution.” *ArXiv* preprint. Available at <https://arxiv.org/abs/1801.04016>.
- Ritchie, Stuart J., Richard Wiseman, and Christopher C. French. 2012. “Failing the Future: Three Unsuccessful Attempts to Replicate Bem’s ‘Retrospective Facilitation of Recall’ Effect.” *PLOS ONE* 7 (3).
- Rohrer, Julia M., Boris Egloff, and Stefan C. Schmukle. 2017. “Probing Birth-Order Effects on Narrow Traits Using Specification-Curve Analysis.” *Psychological Science* 28 (12): 1821–32.
- Salesses, Philip, Katja Schechtner, and César A. Hidalgo. 2013. “The Collaborative Image of The City: Mapping the Inequality of Urban Perception.” *PLOS ONE* 8 (7).
- Salganik, Matthew J., Ian Lundberg, Alexander T. Kindel, Caitlin E. Ahearn, Khaled Al-Ghoneim, Abdullah Almaatouq, Drew M. Altschul, et al. 2020. “Measuring the Predictability of Life Outcomes with a Scientific Mass Collaboration.” *Proceedings of the National Academy of Sciences* 117 (15): 8398–403.

- 
- 
- Silberzahn, Raphael, Eric L. Uhlmann, D. P. Martin, P. Anselmi, F. Aust, E. Awtrey, S. Bahkik, et al. 2018. "Many Analysts, One Dataset: Making Transparent How Variations in Analytic Choices Affect Results." *Advances in Methods and Practices in Psychological Science* 1 (3): 337–56.
- Silberzahn, Raphael, and Eric L. Uhlmann. 2015. "Crowdsourced Research: Many Hands Make Light Work." *Nature* 526:189–91.
- Simonsohn, Uri, Joseph P. Simmons, and Leif D. Nelson. 2015. "Specification Curve: Descriptive and Inferential Statistics on All Reasonable Specifications." Social Science Research Network Working Paper. Available at [https://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2694998](https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2694998).
- Søndergaard, Mikael. 1994. "Research Note: Hofstede's Consequences: A Study of Reviews, Citations, and Replications." *Organization Studies* 15 (3): 447–56.
- Sørensen, Aage B. 1996. "The Structural Basis of Social Inequality." *American Journal of Sociology* 101 (5): 1333–65.
- Stockemer, Daniel, Sebastian Koehler, and Tobias Lentz. 2018. "Data Access, Transparency, and Replication: New Insights from the Political Behavior Literature." *PS: Political Science & Politics* 51 (4): 799–803.
- Uhlmann, Eric L., C. R. Ebersole, C. R. Chartier, T. M. Errington, M. C. Kidwell, C. K. Lai, Randy McCarthy, et al. 2018. "Scientific Utopia: III. Crowdsourcing Science." *PsyArXiv*. Available at <https://psyarxiv.com/vg649>.
- Wicherts, Jelte M., C. L. S. Veldkamp, H. E. M. Augusteijn, M. Bakker, R. C. M. van Aert, and M. A. L. van Assen. 2016. "Degrees of Freedom in Planning, Running, Analyzing, and Reporting Psychological Studies: A Checklist to Avoid p-Hacking." *Frontiers in Psychology* 7:1832.
- Yong, Ed. 2012. "Replication Studies: Bad Copy." *Nature News* 485 (7398): 298.
- Zhao, Yuxiang, and Qinghua Zhu. 2014. "Evaluation on Crowdsourcing Research: Current Status and Future Direction." *Information Systems Frontiers* 16 (3): 417–34.

### Diversity and Inclusion Programs

The American Political Science Association has several major programs aimed at enhancing diversity within the discipline and identifying and aiding students and faculty from under-represented backgrounds in the political science field. These programs include:

#### **Ralph Bunche Summer Institute (RBSI)** (Undergraduate Juniors)

The RBSI Program is an annual five-week program designed to introduce to the world of doctoral study in political science to those undergraduate students from under-represented racial/ethnic groups or those interested in broadening participation in political science and pursuing scholarship on issues affecting underrepresented groups or issues of tribal sovereignty and governance. Application deadline: January of each year. For more information, visit [www.apsanet.org/rbsi](http://www.apsanet.org/rbsi).

#### **APSA Minority Fellows Program (MFP)** (Undergraduate Seniors or MA and PhD students)

(Fall Cycle for seniors and MA Students, Spring Cycle for PhD students) MFP is a fellowship competition for those applying to graduate school, designed to increase the number of individuals from under-represented backgrounds with PhD's in political science. Application deadline: October and March of each year. For more information, visit [www.apsanet.org/mfp](http://www.apsanet.org/mfp).

#### **Minority Student Recruitment Program (MSRP)** (Undergraduates and Departmental members)

The MSRP was created to identify undergraduate students from under-represented backgrounds who are interested in, or show potential for, graduate study and, ultimately, to help further diversify the political science profession. For more information, visit [www.apsanet.org/msrp](http://www.apsanet.org/msrp).

#### **APSA Mentoring Program**

The Mentoring Program connects undergraduate, graduate students, and junior faculty to experienced and senior members of the profession for professional development mentoring. APSA membership is required for mentors. To request a mentor or be a mentor, visit [www.apsanet.org/mentor](http://www.apsanet.org/mentor).

#### **APSA Status Committees**

APSA Status Committees develop and promote agendas and activities concerning the professional development and current status of under-represented communities within the political science discipline. For a listing of all APSA status committees, visit [www.apsanet.org/status-committees](http://www.apsanet.org/status-committees).

For more information on all Diversity and Inclusion Programs, visit us online at [www.apsanet.org/diversityprograms](http://www.apsanet.org/diversityprograms). Please contact Kimberly Mealy, PhD, Senior Director of Diversity and Inclusion Programs with any questions: [kmealy@apsanet.org](mailto:kmealy@apsanet.org).

To contribute to an APSA Fund, such as the Ralph Bunche Endowment Fund or the Hanes Walton Jr. Fund, visit us at [www.apsanet.org/donate](http://www.apsanet.org/donate).