University of Groningen

Crowdsourcing hypothesis tests

Crowdsourcing Hypothesis Tests Collaboration; Albers, Casper

Published in:
Psychological Bulletin

DOI:
[10.1037/bul0000220](https://doi.org/10.1037/bul0000220)

IMPORTANT NOTE: You are advised to consult the publisher's version (publisher's PDF) if you wish to cite from it. Please check the document version below.

Document Version
Publisher's PDF, also known as Version of record

Publication date:
2020

[Link to publication in University of Groningen/UMCG research database](#)

Citation for published version (APA):

Crowdsourcing Hypothesis Tests Collaboration, & Albers, C. (2020). Crowdsourcing hypothesis tests: Making transparent how design choices shape research results. *Psychological Bulletin*, 146(5), 451–479. <https://doi.org/10.1037/bul0000220>

Copyright

Other than for strictly personal use, it is not permitted to download or to forward/distribute the text or part of it without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license (like Creative Commons).

The publication may also be distributed here under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license. More information can be found on the University of Groningen website: <https://www.rug.nl/library/open-access/self-archiving-pure/taverne-amendment>.

Take-down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Downloaded from the University of Groningen/UMCG research database (Pure): <http://www.rug.nl/research/portal>. For technical reasons the number of authors shown on this cover page is limited to 10 maximum.

Psychological Bulletin

Crowdsourcing Hypothesis Tests: Making Transparent How Design Choices Shape Research Results

Justin F. Landy, Miaolei (Liam) Jia, Isabel L. Ding, Domenico Viganola, Warren Tierney, Anna Dreber, Magnus Johannesson, Thomas Pfeiffer, Charles R. Ebersole, Quentin F. Gronau, Alexander Ly, Don van den Bergh, Maarten Marsman, Koen Derks, Eric-Jan Wagenmakers, Andrew Proctor, Daniel M. Bartels, Christopher W. Bauman, William J. Brady, Felix Cheung, Andrei Cimpian, Simone Dohle, M. Brent Donnellan, Adam Hahn, Michael P. Hall, William Jiménez-Leal, David J. Johnson, Richard E. Lucas, Benoît Monin, Andres Montealegre, Elizabeth Mullen, Jun Pang, Jennifer Ray, Diego A. Reinero, Jesse Reynolds, Walter Sowden, Daniel Storage, Runkun Su, Christina M. Tworek, Jay J. Van Bavel, Daniel Walco, Julian Wills, Xiaobing Xu, Kai Chi Yam, Xiaoyu Yang, William A. Cunningham, Martin Schweinsberg, Molly Urwitz, The Crowdsourcing Hypothesis Tests Collaboration, and Eric L. Uhlmann

Online First Publication, January 16, 2020. <http://dx.doi.org/10.1037/bul0000220>

CITATION

Landy, J. F., Jia, M. (L.), Ding, I. L., Viganola, D., Tierney, W., Dreber, A., Johannesson, M., Pfeiffer, T., Ebersole, C. R., Gronau, Q. F., Ly, A., van den Bergh, D., Marsman, M., Derks, K., Wagenmakers, E.-J., Proctor, A., Bartels, D. M., Bauman, C. W., Brady, W. J., Cheung, F., Cimpian, A., Dohle, S., Donnellan, M. B., Hahn, A., Hall, M. P., Jiménez-Leal, W., Johnson, D. J., Lucas, R. E., Monin, B., Montealegre, A., Mullen, E., Pang, J., Ray, J., Reinero, D. A., Reynolds, J., Sowden, W., Storage, D., Su, R., Tworek, C. M., Van Bavel, J. J., Walco, D., Wills, J., Xu, X., Yam, K. C., Yang, X., Cunningham, W. A., Schweinsberg, M., Urwitz, M., The Crowdsourcing Hypothesis Tests Collaboration, & Uhlmann, E. L. (2020, January 16). Crowdsourcing Hypothesis Tests: Making Transparent How Design Choices Shape Research Results. *Psychological Bulletin*. Advance online publication. <http://dx.doi.org/10.1037/bul0000220>

Crowdsourcing Hypothesis Tests: Making Transparent How Design Choices Shape Research Results

Justin F. Landy
Nova Southeastern University

Isabel L. Ding
National University of Singapore

Warren Tierney
University of Limerick

Magnus Johannesson
Stockholm School of Economics

Charles R. Ebersole
University of Virginia

Alexander Ly
University of Amsterdam and Centrum Wiskunde &
Informatica, Amsterdam, the Netherlands

Koen Derks
Nyenrode Business University

Andrew Proctor
Stockholm School of Economics

Christopher W. Bauman
University of California, Irvine

Felix Cheung
University of Hong Kong

Simone Dohle
University of Cologne

Miaolei (Liam) Jia
University of Warwick

Domenico Viganola
George Mason University

Anna Dreber
Stockholm School of Economics and University of Innsbruck

Thomas Pfeiffer
Massey University

Quentin F. Gronau
University of Amsterdam

Don van den Bergh and Maarten Marsman
University of Amsterdam

Eric-Jan Wagenmakers
University of Amsterdam

Daniel M. Bartels
University of Chicago

William J. Brady
New York University

Andrei Cimpian
New York University

M. Brent Donnellan
Michigan State University

Justin F. Landy, Department of Psychology and Neuroscience, Nova Southeastern University; Miaolei (Liam) Jia, Warwick Business School, University of Warwick; Isabel L. Ding, Department of Marketing, National University of Singapore; Domenico Viganola, Department of Engineering, C4I and Cyber Center, George Mason University; Warren Tierney, Kemmy Business School, University of Limerick; Anna Dreber, Department of Economics, Stockholm School of Economics, and Department of Economics, University of Innsbruck; Magnus Johannesson, Department of Economics, Stockholm School of Economics; Thomas Pfeiffer, New Zealand Institute for Advanced Study, Massey University; Charles R. Ebersole, Department of Psychology, University of Virginia; Quentin F. Gronau, Faculty of Social and Behavioural Sciences, University of Amsterdam; Alexander Ly, Faculty of Social and Behavioural Sciences, University of Amsterdam, and Centrum Wiskunde & Informatica, Amsterdam, the Netherlands; Don van den Bergh and Maarten Marsman, Faculty of Social and Behavioural Sciences, University of Amsterdam; Koen Derks, Nyenrode Business University; Eric-Jan Wagenmakers, Faculty of Social and Behavioural Sciences, University of Amsterdam; Andrew Proctor, Department of Economics, Stockholm School of

Economics; Daniel M. Bartels, Department of Marketing, Booth School of Business, University of Chicago; Christopher W. Bauman, Paul Merage School of Business, University of California, Irvine; William J. Brady, Department of Psychology, New York University; Felix Cheung, School of Public Health, University of Hong Kong; Andrei Cimpian, Department of Psychology, New York University; Simone Dohle, Social Cognition Center Cologne, University of Cologne; M. Brent Donnellan, Department of Psychology, Michigan State University; Adam Hahn, Social Cognition Center Cologne, University of Cologne; Michael P. Hall, Department of Psychology, University of Michigan; William Jiménez-Leal, Department of Psychology, University of the Andes; David J. Johnson, Department of Sociology, University of Maryland at College Park; Richard E. Lucas, Department of Psychology, Michigan State University; Benoît Monin, Graduate School of Business and Department of Psychology, Stanford University; Andres Montealegre, Department of Psychology, University of the Andes, and Department of Psychology, Cornell University; Elizabeth Mullen, School of Management, Lucas College, and Graduate School of Business, San José State University; Jun Pang, Department of Marketing, Renmin Business School, Renmin University of China; Jennifer Ray and Diego A. Reinero, Department of

continued

Adam Hahn
University of Cologne

William Jiménez-Leal
University of the Andes

Richard E. Lucas
Michigan State University

Andres Montealegre
University of the Andes and Cornell University

Jun Pang
Renmin University of China

Jesse Reynolds
Stanford University

Daniel Storage
University of Denver

Christina M. Tworek
HarrisX, Rochester, New York

Daniel Walco
New York Yankees, Bronx, New York

Xiaobing Xu
Hainan University

Xiaoyu Yang
Tsinghua University

Martin Schweinsberg
European School of Management and Technology

The Crowdsourcing Hypothesis Tests Collaboration

Michael P. Hall
University of Michigan

David J. Johnson
University of Maryland at College Park

Benoît Monin
Stanford University

Elizabeth Mullen
San José State University

Jennifer Ray and Diego A. Reinero
New York University

Walter Sowden
University of Michigan and Tripler Army Medical Center,
Honolulu, Hawaii

Runkun Su
National University of Singapore

Jay J. Van Bavel
New York University

Julian Wills
New York University

Kai Chi Yam
National University of Singapore

William A. Cunningham
University of Toronto

Molly Urwitz
Stockholm School of Economics

Eric L. Uhlmann
INSEAD

Psychology, New York University; Jesse Reynolds, Department of Psychology, Stanford University; Walter Sowden, Department of Psychology, University of Michigan, and Tripler Army Medical Center, Honolulu, Hawaii; Daniel Storage, Department of Psychology, University of Denver; Runkun Su, NUS Business School, National University of Singapore; Christina M. Tworek, HarrisX, Rochester, New York; Jay J. Van Bavel, Department of Psychology, New York University; Daniel Walco, New York Yankees, Bronx, New York; Julian Wills, Department of Psychology, New York University; Xiaobing Xu, Hainan University; Kai Chi Yam, Department of Management and Organization, National University of Singapore; Xiaoyu Yang, Department of Electronic Engineering, Tsinghua University; William A. Cunningham, Rotman School of Management, University of Toronto; Martin Schweinsberg, Department of Organizational Behavior, European School of Management and Technology; Molly Urwitz, Department of Economics, Stockholm School of Economics; The Crowdsourcing Hypothesis Tests Collaboration; Eric L. Uhlmann, Department of Organizational Behavior, INSEAD.

William J. Brady is now at the Department of Psychology, Yale University.

The first through fourth and last authors developed the research questions, oversaw the project, and contributed equally. The first through third authors oversaw the Main Studies and Replication Studies, and the fourth, sixth, seventh, and eighth authors oversaw the Forecasting Study. The first, fourth,

fifth, eighth, and ninth authors conducted the primary analyses. The 10th through 15th authors conducted the Bayesian analyses. The first and 16th authors conducted the multivariate meta-analysis. The first through ninth and last authors drafted the manuscript. The 17th through 47th authors developed the materials used in the Main Studies and Replication Studies. The 48th author assisted in designing the Forecasting Study. The members of the Crowdsourcing Hypothesis Tests Collaboration lent their expertise as evaluators of study quality and forecasters (see Appendix for details). All authors revised and approved the final manuscript. Data collection was funded by a research grant from INSEAD and a grant from the Jan Wallander and Tom Hedelius Foundation (Svenska Handelsbankens Forskningsstiftelser) to Anna Dreber. Anna Dreber is also funded by the Knut and Alice Wallenberg Foundation (through a Wallenberg Academy Fellows grant) and the Austrian Science Fund (FWF, SFB F63), and Anna Dreber, Magnus Johannesson, and Domenico Viganola are funded by the Jan Wallander and Tom Hedelius Foundation (Svenska Handelsbankens Forskningsstiftelser) and the Swedish Foundation for Humanities and Social Sciences. Thomas Pfeiffer is funded by Marsden Fund Grants 16-UOA-190 and 17-MAU-133.

Correspondence concerning this article should be addressed to Justin F. Landy, Department of Psychology and Neuroscience, Nova Southeastern University, 3301 College Avenue, Fort Lauderdale, FL 33314. E-mail: jlandy@nova.edu

To what extent are research results influenced by subjective decisions that scientists make as they design studies? Fifteen research teams independently designed studies to answer five original research questions related to moral judgments, negotiations, and implicit cognition. Participants from 2 separate large samples (total $N > 15,000$) were then randomly assigned to complete 1 version of each study. Effect sizes varied dramatically across different sets of materials designed to test the same hypothesis: Materials from different teams rendered statistically significant effects in opposite directions for 4 of 5 hypotheses, with the narrowest range in estimates being $d = -0.37$ to $+0.26$. Meta-analysis and a Bayesian perspective on the results revealed overall support for 2 hypotheses and a lack of support for 3 hypotheses. Overall, practically none of the variability in effect sizes was attributable to the skill of the research team in designing materials, whereas considerable variability was attributable to the hypothesis being tested. In a forecasting survey, predictions of other scientists were significantly correlated with study results, both across and within hypotheses. Crowdsourced testing of research hypotheses helps reveal the true consistency of empirical support for a scientific claim.

Public Significance Statement

Research in the social sciences often has implications for public policies as well as individual decisions—for good reason, the robustness of research findings is therefore of widespread interest both inside and outside academia. Yet, even findings that directly replicate—emerging again when the same methodology is repeated—may not always prove conceptually robust to different methodological approaches. The present initiative suggests that crowdsourcing study designs using many research teams can help reveal the conceptual robustness of the effects, better informing the public about the state of the empirical evidence.

Keywords: conceptual replications, crowdsourcing, forecasting, research robustness, scientific transparency

Supplemental materials: <http://dx.doi.org/10.1037/bul0000220.supp>

Scientific theories are meant to be generalizable. They organize findings, ideas, and observations into systems of knowledge that can make predictions across situations and contexts. Theories are more useful when they can explain a wider variety of phenomena. Understanding a theory's scope is critical to successfully applying it. To be generalizable, theories often make use of abstract concepts or conceptual variables to organize their hypothesized relationships. For instance, cognitive dissonance theory, one of the most influential theories in social psychology, states that when individuals have inconsistent cognitions, they will experience psychological distress or discomfort that motivates them to reduce the inconsistency (Festinger, 1957). This theory makes use of conceptual variables to describe its relationships of interest. In particular, *cognitions* refer to any of several types of mental constructs, including attitudes, beliefs, self-concepts, and knowledge that one has engaged in a certain behavior. Reducing inconsistency can take many forms, such as altering one or both of the cognitions to become consistent, or adding new cognitions that resolve the discrepancy. These conceptual variables allow researchers to use the theory to make predictions about many different situations in which people experience inconsistency.

Researchers must operationalize abstract and conceptual variables into concrete terms for empirical testing. For example, to study cognitive dissonance, a researcher might identify two cognitions that could reasonably be brought into conflict with one another (the independent variable). Then, the psychologist could identify a way of resolving the conflict to provide to participants (the dependent variable). Indeed, psychologists have studied cognitive dissonance by measuring attitudes toward a boring task after inducing some participants to lie to the next participant and say the task was exciting (Festinger & Carlsmith, 1959), by measuring preferences toward

appliances after obliging participants to choose between two attractive options to receive as a gift (Brehm, 1956), or by assessing interest in a study group after undergoing an uncomfortable initiation (Aronson & Mills, 1959). Each of these concrete operationalizations widens the understood boundaries of the conceptual variables involved in an effect and thus the generalizability of the effect itself (Schmidt, 2009; Stroebe & Strack, 2014).

Although generalizability is a critical goal of scientific research, the standard model of conducting research creates many challenges for establishing robust generalizability of an effect across contexts. Researchers and/or labs often work in isolation or in small groups, generating their own hypotheses, measures, and operationalizations. These operationalizations represent a small subset of the possible, theoretically justifiable methods that they could have used to test their hypotheses (Baribault et al., 2018; Judd, Westfall, & Kenny, 2012; Monin & Oppenheimer, 2014; Monin, Pizarro, & Beer, 2007; Wells & Windschitl, 1999; Westfall, Judd, & Kenny, 2015). In particular, scientists may use methods that are likely to confirm their preconceptions (McGuire, 1973, 1983; Nickerson, 1998). For example, researchers who theorize that moral judgments are intuitive tend to use simple and emotionally evocative scenarios, whereas researchers who theorize that moral judgments are rooted in reasoning tend to use complex stimuli that pit different values against each other and stimulate deliberation (Monin et al., 2007). Such assumptions may guide which operationalizations are used to test hypotheses and theories, and divergence across operationalizations may then affect which theory is empirically supported.

After one or a few operationalizations and stimulus sets are tested, researchers choose which observations to report to the broader scientific community in academic journals. There is substantial evidence that scientific publishing is biased in favor of

positive or statistically significant findings, leaving negative and null results underreported (Greenwald, 1975; Ioannidis, 2005; Ioannidis & Trikalinos, 2007; Pfeiffer, Bertram, & Ioannidis, 2011; Rosenthal, 1979; Schimmack, 2012; Simonsohn, Nelson, & Simmons, 2014). Null results are important for understanding generalizability because they provide insights about where the boundaries of a theory lie; nonetheless, the scientific community may be left largely unaware of them due to biases in publishing (LeBel, McCarthy, Earp, Elson, & Vanpaemel, 2018; Zwaan, Etz, Lucas, & Donnellan, 2018).

After initial observations are reported, other researchers may conduct follow-up research. These follow-ups have the potential to increase understanding of generalizability by inspiring new operationalizations and instantiations of effects and theories. Still, scientific culture and professional advancement often privilege novelty over increased certainty and incremental refinement (Everett & Earp, 2015; Giner-Sorolla, 2012; Nosek, Spies, & Motyl, 2012), which may disincentivize researchers from conducting tests of previously published ideas in favor of pursuing new ideas and theories (Makel, Plucker, & Hegarty, 2012). Although scientific culture has been changing with respect to valuations of replications, particularly in psychology, these changes have been more focused on direct replications (testing the same idea with the same materials and methodology; Alogna et al., 2014; Ebersole et al., 2016; Klein et al., 2014; Open Science Collaboration, 2015; Pashler & Harris, 2012; Simons, 2014) than on conceptual replications (testing established ideas with a new approach; Crandall & Sherman, 2016; Finkel, Eastwick, & Reis, 2015, 2017). Furthermore, failed conceptual replications are far more susceptible to alternative explanations based on methodological differences than are direct replications, and as a consequence may be left unpublished or dismissed by original researchers and other scientists (Baribault et al., 2018; Doyen, Klein, Simons, & Cleeremans, 2014; Earp, in press; Hendrick, 1990; Schmidt, 2009; Simons, 2014). Taken together, these forces within the standard model of conducting psychological research may impede tests of generalizability of scientific theories and phenomena. The standard model may thus stunt theory development by limiting contributions to the literature to ones based on a relatively small subset of operationalizations, and to unrealistically positive results.

The Current Research

To address these challenges, we introduce a crowdsourced approach to hypothesis testing. In the crowdsourcing initiative reported here, up to 13 research teams (of a total of 15 teams) independently created stimuli to address the same five research questions, while fully blind to one another's approaches, and to the original methods and the direction of the original results. The original hypotheses, which were all unpublished at the time the project began, dealt with topics including moral judgment, negotiations, and implicit cognition. Large samples of research participants were then randomly assigned to different teams' versions of the same study, with a commitment to publish the results from all study designs as a fundamental component of the project. The analyses were also preregistered, which has been argued to reduce bias (Nosek, Ebersole, DeHaven, & Mellor, 2018; Wagenmakers, Wetzels, Borsboom, van der Maas, & Kievit, 2012; Wicherts et al., 2016), although a causal effect remains to be empirically

demonstrated. Comparisons of the estimated effect sizes associated with the same hypothesis across the studies created by the different teams reveal the extent to which the empirical results are contingent on the decisions scientists make as they design their study. Aggregating results across teams via meta-analysis, taking into account both average effects and variability across teams, provides a systematic assessment of the relative strength of support for each hypothesis.

There are a number of potential benefits to a crowdsourcing approach to hypothesis testing. Crowdsourcing the operationalization of research ideas makes transparent the true consistency of support for an empirical prediction, and provides a more stringent test of robustness than employing a narrow set of stimuli (Monin & Oppenheimer, 2014), directly replicating multiple independent and dependent variables that have been used previously (Caruso, Shapira, & Landy, 2017), or even the innovative approach of radically randomizing features of the same basic experimental design (e.g., symbols, colors, and presentation speeds in a cognitive priming paradigm; Baribault et al., 2018). Rather than varying features of the same basic design to address concerns about stimulus sampling (Baribault et al., 2018), we had different researchers design distinct studies to test the same research questions, providing an arguably wider-ranging test of the conceptual robustness of each original finding. The extent to which divergent approaches produce different results is further revealed. Uniquely, the conceptual replications are developed by independent research teams with no prior knowledge of the original authors' method or results to bias them (Silberzahn et al., 2018; Silberzahn & Uhlmann, 2015), unlike in the usual practice of science, in which conceptual replications are conducted after the dissemination of the original results. Materials designers also did not know the direction of the original hypotheses and results, but were rather provided with a nondirectional version of each research question (see below). This was to prevent materials designers from constructing materials aimed at confirming a directional hypothesis, while not giving alternative directional hypotheses a chance (Monin et al., 2007). In other words, we believe that we were more concerned with answering the research questions that drove the five original, unpublished studies than we were with confirming their results.

Because all crowdsourced conceptual replications were preregistered and reported, this approach is also free of reporting bias, unlike traditional conceptual replications where null effects may be attributed to departures from the original methodology and therefore left unpublished. Moreover, because participants from the same large sample are randomly assigned to different conceptual replications, discrepant results, including failed replications, cannot be attributed to differences in the populations being sampled (McShane, Tackett, Böckenholt, & Gelman, 2019; Tiokhin, Hackman, Munira, Jesmin, & Hruschka, 2019; Van Bavel, Mende-Siedlecki, Brady, & Reinero, 2016). Heterogeneity in results above-and-beyond what would be expected based on sampling error can confidently be attributed to design choices.

In the present initiative, we also recruited a second large sample and repeated our initial studies with the same methodologies and materials. This effort is, to our knowledge, the first time an entire crowdsourced set of studies has itself been directly replicated. Doing so allowed us to simultaneously take into account both conceptual and direct replications when assessing the strength of evidence for each finding. Altogether, we provide a new frame-

work for determining the generalizability and context-dependency of new findings, with the goal of identifying more deeply robust phenomena, which we believe may hold utility for select research questions in the future. In the Discussion, we elaborate at greater length on when crowdsourcing hypothesis tests is likely to prove most (and least) useful.

We additionally examine whether scientists are able to predict *a priori* how design choices impact research results. Prior work has demonstrated that researchers can anticipate whether a published result will independently replicate based on the research report alone (Camerer et al., 2016; Dreber et al., 2015), and predict the effects of performance interventions starting only from a few benchmark effects and the materials for the additional treatments (DellaVigna & Pope, 2018a, 2018b). Other forecasting studies with scientists have returned more mixed results (Coffman & Niehaus, 2014; Dunaway, Edmonds, & Manley, 2013; Groh, Krishnan, McKenzie, & Vishwanath, 2016; Sanders, Mitchell, & Chonaire, 2015). We therefore conducted a forecasting survey asking an independent crowd of scientists to attempt to predict the results of each study based solely on its sample size, methodology, and materials. Notably, all prior work has examined forecasting accuracy *across* different hypotheses that vary in their truth value and alignment with empirical reality. In contrast, we assessed whether scientists are accurate in their beliefs about the outcomes of *different* experiments designed to test the *same* research question. Scientists' intuitions about the impact of researcher choices may or may not map onto the actual downstream consequences.

Method

Main Studies and Replication Studies

In two separate data collection efforts (the initial investigations ["Main Studies"] and direct replications ["Replication Studies"]), we randomly assigned participants to different sets of study materials designed independently by up to 13 teams of researchers to test the same five research questions. There were 15 teams of materials designers in total, from which up to 13 teams designed materials for each research question (i.e., not all teams made materials to test all five original hypotheses). The five research questions were gathered by emailing colleagues conducting research in the area of moral judgment and asking if they had initial evidence for an effect that they would like to volunteer for crowdsourced testing by other research groups. In three cases (Hypotheses 1, 3, and 4), project coordinators volunteered an effect from their research program that fit these criteria, and in two cases, members of other teams volunteered an effect (Hypotheses 2 and 5). In the present research, we examined the overall degree of support for each hypothesis, and also quantified the heterogeneity across different sets of study materials. To our knowledge, this instance is the first time a large-scale metascientific project has itself been directly replicated in full with a new sample. All materials, data, and analysis scripts from this project are publicly available at <https://osf.io/9jzy4>.

Target Hypotheses

We identified five directional hypotheses in the areas of moral judgment, negotiation, and implicit cognition, each of which had been

supported by one then-unpublished study.¹ Table 1 shows the directional hypotheses, as well as the nondirectional forms in which they were presented to materials designers. Below we elaborate briefly on the theoretical basis for each research question.

Hypothesis 1: Awareness of automatic prejudice. Influential dual-process theories of intergroup attitudes propose that individuals have both explicit, consciously endorsed attitudes toward negatively stereotyped groups, and also implicit ones that they may not endorse (Dovidio, Kawakami, & Gaertner, 2002; Fazio, Jackson, Dunton, & Williams, 1995; Gawronski & Bodenhausen, 2006; Greenwald & Banaji, 1995; Wilson, Lindsey, & Schooler, 2000). Rather than in propositional logic, these implicit attitudes are based in simple associations (e.g., Black-Criminal, Female-Weak), that are conditioned by the cultural environment (Gawronski & Bodenhausen, 2006; Uhlmann, Poehlman, & Nosek, 2012). As a result, even consciously egalitarian individuals often harbor prejudiced associations that may leak out and affect their judgments and behaviors without them realizing it. Low correspondence between self-reported and implicit measures of intergroup attitudes has been interpreted as indicating a lack of introspective access into the latter (Banaji, 2001; Greenwald & Banaji, 1995). Nonetheless, people could potentially be aware of their spontaneous affective reactions without endorsing them. Indeed, Hahn and colleagues (Hahn, Judd, Hirsh, & Blair, 2014; Hahn & Gawronski, 2019) provide evidence that people can accurately predict their performance on Implicit Association Tests of associations with social groups (Greenwald, McGhee, & Schwartz, 1998). Uhlmann and Cunningham (2000) constructed questionnaire items examining whether individuals directly self-report negative gut feelings toward minorities. Representative items include "Although I don't necessarily agree with them, I sometimes have prejudiced feelings (like gut reactions or spontaneous thoughts) that I don't feel I can prevent" and "At times stereotypical thoughts about minorities coming into my head without my necessarily intending them to." In the original research, approximately three-quarters of undergraduates agreed with such statements, and overall endorsement was confirmed by mean responses statistically significantly above the neutral scale midpoint of four (1 = *strongly disagree*, 4 = *neutral*, 7 = *strongly agree*). As the Uhlmann and Cunningham (2000) investigations were never published, the present initiative crowdsourced the question of whether people self-report automatic intergroup prejudices, assigning a dozen independent research teams to create their own awareness measures. Specifically, we examined whether the majority of people, without further prompting or consciousness-raising, agree on questionnaire measures that they harbor such automatic biases toward stigmatized groups.

Hypothesis 2: Extreme offers reduce trust. Negotiators are routinely advised to make extreme first offers to benefit from the anchoring effect (Tversky & Kahneman, 1974). When sellers make extreme first offers, final prices tend to be high; in contrast, when buyers make extreme first offers, final prices tend to be low (Ames & Mason, 2015; Galinsky, Leonardelli, Okhuysen, & Mussweiler, 2005; Galinsky & Mussweiler, 2001). Evidence suggests this effect is robust across cultures, issues, and power positions (Gunia, Swaab, Sivathan, & Galinsky, 2013). Yet, more recent research has examined the

¹ The original study supporting Hypothesis 4 has since been published as a supplemental study in Landy et al. (2017).

Table 1

Directional and Nondirectional Formulations of the Five Hypotheses

Hypothesis 1	
Directional:	People explicitly self-report an awareness of harboring negative automatic associations with members of negatively stereotyped social groups.
Nondirectional:	When directly asked, do people explicitly self-report an awareness of harboring negative automatic associations with members of negatively stereotyped social groups?
Hypothesis 2	
Directional:	Negotiators who make extreme first offers are trusted less, relative to negotiators who make moderate first offers.
Nondirectional:	Are negotiators who make extreme first offers trusted more, less, or the same relative to negotiators who make moderate first offers?
Hypothesis 3	
Directional:	A person continuing to work despite having no material/financial need to work has beneficial effects on moral judgments of that individual.
Nondirectional:	What are the effects of continuing to work despite having no material/financial need to work on moral judgments of that individual—beneficial, detrimental, or no effect?
Hypothesis 4	
Directional:	Part of why people are opposed to the use of performance enhancing drugs in sports is because they are “against the rules.” But, whether the performance enhancer is against the rules established by a proximal authority (e.g., the league) contributes more to this judgment than whether it is against the law.
Nondirectional:	Part of why people are opposed to the use of performance enhancing drugs in sports is because they are “against the rules.” But which contributes more to this judgment—whether the performance enhancer is against the law, or whether it is against the rules established by a more proximal authority (e.g., the league)?
Hypothesis 5	
Directional:	The tendency to make deontological (as opposed to utilitarian) judgments is positively related to personal happiness.
Nondirectional:	Is a utilitarian vs. deontological moral orientation related to personal happiness?

conditions under which this advice might not be accurate (Loschelder, Swaab, Trötschel, & Galinsky, 2014; Loschelder, Trötschel, Swaab, Friese, & Galinsky, 2016; Maaravi & Levy, 2017). The present Hypothesis 2 explores one mechanism for why extreme first offers might backfire in negotiations with multiple issues. Specifically, extreme first offers may interfere with value creation processes such as trust building and information exchange. Building on previous research that showed that extreme first offers can cause offense and even impasses (Schweinsberg, Ku, Wang, & Pillutla, 2012), Schweinsberg (2013) examined the specific hypothesis that extreme first offers lower trust in the counterpart. Ultimately, this line of research may show that extreme first offers can help negotiators claim a larger percentage of the bargaining zone for themselves, but that extreme first offers also shrink the overall size of the bargaining zone by reducing information exchange and trust. Thus, extreme first offers might help negotiators claim a larger percentage of a smaller bargaining zone, making them ultimately worse off. Negotiators might be blind to this extreme first offer disadvantage because their salient comparison is between value they claimed versus value claimed by their counterpart, and not the more relevant but counterfactual comparison between value they claimed from an extreme offer versus value they could have claimed from a more moderate first offer. The present research focuses on just one part of this argument, providing crowdsourced tests of the prediction that “negotiators who make extreme first offers are trusted less, relative to negotiators who make moderate first offers.”

Hypothesis 3: Moral praise for needless work. It is easy to find anecdotal examples in which individuals received moral praise for continuing to work despite coming into sudden wealth and no longer needing to earn a salary (Belsie, 2011). In scenario studies based on such real life cases, Americans positively evaluate

the moral character of individuals with working class occupations (e.g., potato peeler in a restaurant kitchen) who continue their employment after winning the lottery (Poehlman, 2007; Uhlmann, Poehlman, & Bargh, 2009). A number of sources for such moral intuitions are plausible, among these a tendency to value work contributions that parallels general disapproval of shirkers and noncontributors (Jordan, Hoffman, Bloom, & Rand, 2016), use of work behavior as a signal of underlying traits (Uhlmann, Pizarro, & Diermeier, 2015), the influence of the Protestant work ethic in some cultures (Uhlmann & Sanchez-Burks, 2014), and postmaterialist value systems in which work is pursued for meaning and fulfillment rather than as an economic necessity (Inglehart, 1997; Inglehart & Welzel, 2005). A separate project to this one examines the extent to which these and other work morality effects directly replicate across different national cultures (Tierney, Ebersole, Hardy, Chapman, et al., 2019). Of interest to the present initiative is how conceptually robust the findings are to alternative study designs. We therefore crowdsourced the research question of whether working in the absence of material need elicits moral praise, limiting our samples to U.S.-based participants, the group originally theorized to exhibit these effects.

Hypothesis 4: Proximal authorities drive legitimacy of performance enhancers. People in the United States express widespread normative opposition to the use of Performance-Enhancing Drugs (PEDs), especially among competitive athletes, but it is not clear what underpins these judgments. While most studies of opposition to PEDs have examined perceptions of fairness (e.g., Dodge, Williams, Marzell, & Turrissi, 2012; Fitz, Nadler, Manogaran, Chong, & Reiner, 2014; Scheske & Schnall, 2012), some research also suggests that the sheer fact that PEDs are prohibited also contributes to opposition toward them (Sattler, Forlini, Racine, & Sauer, 2013).

This distinction between fairness concerns and explicit rules roughly parallels the insight from Social Domain Theory (Turiel, 1983, 2002) that acts can be wrong in at least two qualitatively different ways: moral offenses violate universal moral standards like fairness, whereas conventional offenses violate consensually accepted norms or the dictates of legitimate authorities. Landy, Walco, and Bartels (2017) investigated whether opposition to PED use exhibits properties of conventional offenses by manipulating whether or not an athlete's use of PEDs "violates the law and the rules of his [competition] circuit" (Study 2 of the original report), and found that this manipulation significantly affected people's judgments of how wrong it was for the athlete to use PEDs. A follow-up study (Supplemental Study 1 of the original report) found that PED use was considered more wrong when it violated a dictate of a legitimate proximal authority (the competition circuit) than when it violated the law. An additional study replicated this finding (Study 12 of the original report), but a further study did not (Study 13 of the original report), so it is unclear whether proximal authority or legal authority contributes more to opposition to PED use. Because all of these studies were unpublished at the beginning of this project, we applied our crowdsourcing methodology to obtain a more definitive answer to this question.

Hypothesis 5: The tendency to make deontological judgments is positively correlated with happiness. To bridge the normative-descriptive divide between the fields of philosophical ethics (how should people morally behave) and moral psychology (how and why do people morally behave) cognitive science must map out how variation in moral cognitions are systematically related to variances in outcomes related to human flourishing. The goal of this original research was to contribute to this endeavor by examining how the tendency to make utilitarian versus deontological moral judgments (Bentham, 1970/1823; Kahane, 2015; Kant, 1993/1785; Mill, 2004/1861) relates to personal happiness and well-being (Kahneman, Diener, & Schwarz, 1999; Ryff, 1989; Waterman, 1993). The idea that happiness and morality are tightly intertwined has a long history in philosophy (see, e.g., Annas, 1993; Aristotle, 340 BCE/2002; Foot, 2001; Kraut, 1979), and recent empirical work suggests that people consider moral goodness to be an element of what "happiness" consists of (Phillips, de Freitas, Mott, Gruber, & Knobe, 2017; Phillips, Nyholm, & Liao, 2014). However, prior work has not examined the relationship (if any) between specific *moral orientations* and happiness.

Hypothesis 5 posits that people who are more inclined to base their moral judgments on the violation of rules, duties, and obligations (deontological judgments) versus material outcomes (utilitarian judgments) are also more likely to experience happiness in their lives. This prediction is based on philosophical and scientific evidence that has demonstrated shared psychological and neurological mechanisms between these dimensions (e.g., Everett, Pizarro, & Crockett, 2016; Greene, 2013; Lieberman, 2013; Phillips et al., 2017; Singer, 2005). To test this hypothesis, Sowden and Hall (2015) asked participants to judge several morally questionable behaviors that pitted utilitarian and deontological considerations against one another (Greene, Sommerville, Nystrom, Darley, & Cohen, 2001) and compared an index of those judgments with how they responded to measures of subjective well-being (Diener, Emmons, Larsen, & Griffin, 1985; Watson, Clark, & Tellegen, 1988) and Eudaimonic happiness (Waterman et al., 2010). The crowdsourced project posed the research question to independent researchers,

who separately designed studies relating moral judgments to individual happiness.

Method

Materials

A subset of the project coordinators (Landy, Jia, Ding, Uhlmann) recruited 15 teams of researchers through their professional networks to independently design materials to test each hypothesis. Of these 15 teams, four included the researchers who developed the original materials for at least one of the five hypotheses. Teams ranged in size from one researcher to five, and members ranged in experience from graduate students to full professors. We opted not to standardize team size because research teams vary greatly in size in the natural practice of science (Wuchty, Jones, & Uzzi, 2007a, 2007b). All studies were required to be designed to be administered in an online survey. Note that recruiting through our professional networks would, if anything, be expected to bias our results toward homogeneity and consistency between materials designers. Likewise, the restriction to using only brief, online questionnaires rather than behavioral measures, video stimuli, or elaborate laboratory experiments with a cover story and research confederates, also artificially constrains variability in study designs. Yet, as we detail below, we still observed remarkable heterogeneity in results.

To avoid biasing their designs, materials designers were provided with the nondirectional versions of the five hypotheses presented in Table 1, and developed materials to test each hypothesis independently of the other teams. The team of Xu and Yang designed materials only for Hypotheses 1, 2, and 5, and the team of Cimpian, Tworek, and Storage designed materials only for Hypotheses 3 and 4. We also included the original materials from the unpublished studies that initially supported each hypothesis and conducted direct replications with them; the teams of Uhlmann, Schweinsberg, and Uhlmann and Cunningham only contributed these original materials. The original materials for Hypothesis 5 were developed by the team of Sowden and Hall, but were much longer than any other materials set, so this team also developed a shorter set of materials for Hypothesis 5 and data were collected using both versions. In all, 64 sets of materials, including the five sets of original materials, were created through this crowdsourced process. The materials and analyses for both studies were preregistered at <https://osf.io/9jzy4/> (see also Supplement 1, as well as Supplement 2 for deviations from the preregistered analyses, in the online supplemental materials).

Participants

In total, 8,080 participants located in the United States began the Main Studies on Amazon Mechanical Turk (MTurk; Chandler, Mueller, & Paolacci, 2014; Chandler, Paolacci, & Mueller, 2013); of these, 7,500 completed the entire study. In accordance with our preregistered stopping rule (see <https://osf.io/avnuc/>), we ceased data collection after $N = 7,500$ participants finished the survey. In the Replication Studies, 7,500 English-speaking adult participants located in the United States were recruited via PureProfile, a survey firm—we employed this different sampling method for the Replication because,

in the Main Studies, we had already essentially exhausted the number of Mechanical Turk participants that a typical lab samples (see Stewart et al., 2015). In both data collection efforts, responses from participants who completed their assigned materials for one or more hypotheses but did not complete all of the assigned materials in their entirety were retained, resulting in slightly different sample sizes across the five hypotheses (Main Studies: Hypothesis 1 $N = 7,175$; Hypothesis 2 $N = 7,160$; Hypothesis 3 $N = 7,146$; Hypothesis 4 $N = 7,158$; Hypothesis 5 $N = 7,758$; Replication Studies: Hypothesis 1 $N = 7,586$; Hypothesis 2 $N = 7,631$; Hypothesis 3 $N = 7,568$; Hypothesis 4 $N = 7,576$; Hypothesis 5 $N = 8,231$). On a per-cell basis, there were approximately 300 participants for Hypotheses 1–4, and 600 participants for Hypothesis 5 (which was tested using a Pearson correlation, rather than a comparison between experimental groups).

Procedure

In the Main Studies, participants were randomly assigned to one set of materials for each of the five hypotheses, and, for designs with multiple conditions, one condition per hypothesis. The order in which the five sets of materials were presented was randomized for each participant. After responding to all five sets of materials, participants completed a demographics questionnaire including questions about their age, gender, and other characteristics. Additionally, a separate subsample of participants was randomly assigned to only complete the full original materials for Hypothesis 5, because of their length. The materials designed by the team of Jiménez-Leal and Montealegre to test Hypothesis 4 were run separately approximately two months after the rest of Main Studies were run, because we discovered that, owing to a coding error, one of the two conditions from these materials was not presented to participants in the original run (new data were therefore collected for both conditions of this design). The procedure for the Replication Studies was essentially identical to that of the Main Studies; the only modifications were fixing the aforementioned condition missing from Hypothesis 4, and preregistering some exploratory analyses conducted on the data from the Main Studies (see Supplement 2 in the online supplemental materials), this time as confirmatory tests (see <https://osf.io/8s69w/>).

Forecasting Study

The online Forecasting Study was open to any scientist and had two purposes. First, it tested the extent to which researchers ($N = 141$) were able to predict the results of the Main Studies and Replication Studies, in terms of the standardized effect size that would be obtained from each set of materials, and also with regard to statistical significance (the likelihood that a p value below .05 would be found). Second, it determined how independent reviewers evaluate each set of materials based on whether it provides an adequate test of the original hypothesis. Variability across different study versions is far more meaningful if they provide valid tests of the original research idea. We placed half of the forecasters at random into a monetarily incentivized version of the survey; potential payoff ranged between \$0 and \$60, meaning financial incentives were present in the treatment condition but not strong. Further methodological details for the forecasting survey can be found in Supplements 3 and 5 in the online supplemental materials, and the preregistration can be found at <https://osf.io/9jzy4/>.

Results

Main Studies and Replication Studies

Given our key theoretical question regarding heterogeneity in estimates, as well as large sample sizes that might render even small and theoretically uninteresting differences statistically significant, our primary focus is on dispersion in effect sizes across different study designs. Yet, because the $p < .05$ threshold is widely used as the lower bound criterion for concluding the presence of an effect, we likewise examined patterns of statistical significance, both at the level of individual designs and aggregated across them. This reliance on both effect sizes and statistical significance levels to quantify the project results was preregistered in advance. Because of the potential issues associated with relying on statistical significance to draw conclusions, we report the results of null hypothesis significance tests in Supplement 9 in the online supplemental materials, and focus here on the analyses of effect sizes.

Meta-analytic statistics. To examine the support for each hypothesis, as well as the variation across study designs for each of them, we computed effect size estimates for the results from each of the 64 sets of materials. The diversity in effect size estimates from different study designs created to test the same theoretical ideas constitute the primary output of this project. For Hypotheses 1–4, the effect sizes were independent-groups Cohen's d s, and for Hypothesis 5, they were Pearson r s. Effect size estimates and sampling variances were calculated via bootstrapping, using the *bootES* package for R (Kirby & Gerlanc, 2013),² then combined in random-effects meta-analyses using the *metafor* package (Viechtbauer, 2010), to obtain an overall estimate for the size of each hypothesized effect.³ This model treats each observed effect size y_i as a function of the average true effect size μ ,

² Materials designed by the team of Donnellan, Lucas, Cheung, and Johnson for Hypotheses 2, 3, and 4 employed within-subjects designs, whereas the other materials for these hypotheses employed one-sample or between-subjects designs. To ensure that all effect sizes were comparable in the meta-analyses, the repeated-measures d s for the within-subjects designs were converted to independent-groups d s (see Morris & DeShon, 2002). *bootES* does not have a feature to convert between effect size metrics, so custom bootstrapping code was used (see <https://osf.io/avnucl/>). This custom code returns the same effect size estimates and variance terms for the repeated-measures d s as *bootES*, and converts the repeated-measures d s to independent-groups d s according to Equation 11 in Morris and DeShon (2002).

³ Fixed-effects models showed similar estimated mean effect sizes. In the Main Studies, the point estimate was not statistically significant for Hypothesis 1, $p = .093$, and was statistically significant for Hypothesis 4, $p < .001$. In the Replication Studies, the estimated effect sizes were again similar when fixed-effects models were used, but the point estimates for Hypotheses 1 and 4 were statistically significant, $p < .001$. Yet, fixed-effects models are not generally recommended when meta-analyzing studies with different methods (Borenstein, Hedges, Higgins, & Rothstein, 2010; Hunter & Schmidt, 2000), so we focus on the random-effects models.

between-study variability, $u_i \sim N(0, \tau^2)$, and sampling error, $e_i \sim N(0, \nu_i)$; see Viechtbauer, 2010)⁴:

$$y_i = \mu + u_i + e_i$$

The heterogeneity among effect sizes (τ^2) was estimated using Restricted Maximum Likelihood Estimation. Positive effect sizes indicate results consistent with the original, unpublished findings, whereas negative effect sizes indicate results in the opposite direction. Figure 1 presents forest plots of the observed effect sizes in these analyses. For ease of comparison across the five figures, the Pearson r effect sizes for Hypothesis 5 have been converted to Cohen's d s (Rosenthal & DiMatteo, 2001). The top panel of each figure presents observed effect sizes and the estimated mean effect size from the Main Studies, and the middle panel presents observed effect sizes and the estimated mean effect size from the Replication Studies. Beneath these panels, the estimated mean effect size for each hypothesis, computed by combining all individual effect sizes in the Main Studies and Replication Studies ($k = 26$ for Hypotheses 1, 2, 3, and 5; $k = 24$ for Hypothesis 4) is presented. The bottom panel presents effect sizes computed by meta-analytically combining the Main Studies' and Replication Studies' effect sizes for each set of materials (i.e., this panel presents the results of 12 or 13 meta-analyses, each with $k = 2$ studies).⁵

In the Main Studies, these analyses showed a statistically significant aggregated effect in the expected direction for Hypotheses 2, 3, and 5 (estimated mean effect sizes: $d = 1.04$, 95% CI [0.61, 1.47], $p < .001$; $d = 0.33$, 95% CI [0.17, 0.50], $p < .001$; $r = .06$, 95% CI [0.01, 0.11], $p = .010$), and no statistically significant aggregated effect as expected under Hypotheses 1 and 4 ($d = 0.07$, 95% CI [-0.22, 0.37], $p = .623$; $d = 0.07$, 95% CI [-0.05, 0.20], $p = .269$). Note that in the case of Hypothesis 5, the aggregated estimate was very small, and the threshold for statistical significance may only have been crossed due to the large sample and the resultant high power to detect even trivially small effects. In the Replication Studies, the patterns of results were similar, though the estimated mean effect sizes tended to be somewhat smaller, overall. Hypotheses 2 and 3 ($d = 0.60$, 95% CI [0.32, 0.88], $p < .001$; $d = 0.24$, 95% CI [0.11, 0.38], $p < .001$) were associated with a statistically significant effect in the expected direction. Hypothesis 5 did not receive statistically significant overall support in the Replication Studies ($r = .03$, 95% CI [-0.04, .09], $p = .417$), though the estimated mean effect size was not meaningfully different than in the Main Studies. Consistent with the Main Studies, Hypotheses 1 and 4 were again not supported ($d = -0.07$, 95% CI [-0.33, 0.19], $p = .588$; $d = 0.03$, 95% CI [-0.05, 0.19], $p = .465$). Overall, then, the meta-analytic results were largely consistent across the Main Studies and the Replication Studies, reflecting overall support for Hypotheses 2 and 3, and an overall lack of support for Hypotheses 1, 4, and 5. Similar results were found when relying on null hypothesis significance testing (see Supplement 9 in the online supplemental materials).

Just as importantly, inspection of the forest plots suggests substantial variation among effect sizes, even within the same hypothesis. We assessed this more formally by examining the Q , I^2 , and τ^2 statistics in each meta-analysis (Huedo-Medina, Sánchez-Meca, Marín-Martínez, & Botella, 2006). The Q statistic is a test for heterogeneity—a significant Q statistic means that heterogeneity in true effects can be expected. Because all participants in each

study were drawn from the same large online sample and randomly assigned to conditions, it is unlikely that heterogeneity can be attributed to hidden moderators (e.g., different populations being sampled, different study environments, etc., see Van Bavel et al., 2016), and thus is likely attributable to differences in the materials. The I^2 statistic quantifies the percentage of variance among effect sizes attributable to heterogeneity, rather than sampling variance. By convention, I^2 values of 25%, 50%, and 75% indicate low, moderate, and high levels of unexplained heterogeneity, respectively (Higgins, Thompson, Deeks, & Altman, 2003). Yet, Q and I^2 are also sensitive to sample size; large samples tend to produce large and significant Q statistics and large I^2 values. Therefore, we also report the τ^2 statistic as an absolute measure of the amount of heterogeneity in our data. The τ^2 statistic is an estimate of the variance of true effect sizes (Borenstein, Hedges, Higgins, & Rothstein, 2009). All five hypotheses showed statistically significant and high levels of heterogeneity in the Main Study and the Replication (see Table 2). In the Main Study, only about 1%, 2%, 6%, 12%, and 24% of the variance across the effect sizes for Hypotheses 1, 2, 3, 4, and 5, respectively, can be attributed to chance variation. Similarly, in the Replication, we would only expect to observe about 1%, 3%, 9%, 22%, and 14% of the variance across the effect sizes for Hypotheses 1, 2, 3, 4, and 5, respectively, by chance. Most of the observed variance across effect sizes in both studies is unexplained heterogeneity. Moreover, the τ^2 statistics are rather large, relative to the estimated mean effect sizes, suggesting that these large I^2 values are not simply due to our large effect sizes resulting in low sampling variance—there are meaningful levels of absolute heterogeneity in our data. One can also see this pattern simply by visually inspecting the forest plots (Figure 1), which show considerable dispersion among effect sizes.

Explaining heterogeneity in effect sizes. We therefore sought to explain this observed heterogeneity. First, we computed intraclass correlation coefficients (ICCs) predicting observed effect sizes from the hypothesis they tested, and the team that designed the materials (see Klein et al., 2014, for a similar analysis). To compare across all observed effect sizes, the Pearson r s from Hypothesis 5 were converted into Cohen's d s (Rosenthal & DiMatteo, 2001), as above. In the Main Studies, the hypothesis being tested was moderately predictive of observed effect sizes, $ICC = .40$, 95% CI [.15, .86], whereas team did not explain

⁴ This analytic approach is not ideal, because it ignores the multivariate nature of the data: each hypothesis can be thought of as a separate outcome variable. It also ignores the multilevel nature of the data (designs are nested within hypotheses), and individual-level correlations across designs resulting from the fact that each participant completed up to five different study designs. We therefore also ran a one-stage multivariate meta-analysis on our individual participant data to model these aspects of the data. The results are very similar to the reported univariate meta-analyses, and this approach has its own disadvantages, particularly that analysis of heterogeneity jointly across outcomes jointly is complicated by the non-nested participant design (see Supplement 8 in the online supplemental materials). Therefore, we focus here on the more familiar analytic approach.

⁵ When meta-analytically combining the Main Studies' and Replication Studies' effect sizes for each individual set of materials, we employed fixed-effects models, unlike in the rest of our meta-analytic models. This is because the two effect sizes being combined come from studies with identical materials and methods, so they should, in principle, be measuring the same true population effect size.

Figure 1. (a) Forest plot of observed effect sizes (independent-groups Cohen's *ds*) for Hypothesis 1. The research question was "When directly asked, do people explicitly self-report an awareness of harboring negative automatic associations with members of negatively stereotyped social groups?" (b) Forest plot of observed effect sizes (independent-groups Cohen's *ds*) for Hypothesis 2. The research question was "Are negotiators who make extreme first offers trusted more, less, or the same relative to negotiators who make moderate first offers?" (c) Forest plot of observed effect sizes (independent-groups Cohen's *ds*) for Hypothesis 3. The research question was "What are the effects of continuing to work despite having no material/financial need to work on moral judgments of that individual—beneficial, detrimental, or no effect?" (d) Forest plot of observed effect sizes (independent-groups Cohen's *ds*) for Hypothesis 4. The research question was "Part of why people are opposed to the use of performance enhancing drugs in sports is because they are 'against the rules.' But which contributes more to this judgment - whether the performance enhancer is against the law, or whether it is against the rules established by a more proximal authority (e.g., the league)?" (e) Forest plot of observed effect sizes (converted to Cohen's *ds*, for comparison to other hypotheses) for Hypothesis 5. The research question was "Is a utilitarian vs. deontological moral orientation related to personal happiness?"

statistically significant variance, $ICC = -.13$, 95% CI $[-.23, .09]$. The negative ICC for team indicates that between-team variance is lower than within-team variance. This means that which team designed a set of materials had no predictive relationship with the

observed effect size (see Bartko, 1976). In other words, some teams were not better than others at designing study materials that produced large effect sizes across hypotheses. We followed up this analysis with a random-effects metaregression, predicting effect

Figure 1. (continued)

sizes from hypothesis and team, with the median hypothesis (Hypothesis 5, in the Main Studies) and the median team (Sowden & Hall) as the reference levels. Hypothesis 2 produced statistically significantly larger effect sizes than the median hypothesis, $\beta = 0.85$, 95% CI [0.47, 1.23], $p < .001$, but, consistent with the analysis above, no team produced statistically significantly larger or smaller effect sizes than the median team, $ps > .086$. Moreover, after accounting for both hypothesis and team, there was still substantial and statistically significant residual heterogeneity across effect sizes, $Q(44) = 1291.64$, $p < .001$, $I^2 = 97.39\%$, 95% CI [96.22, 98.40], $\tau^2 = 0.24$, 95% CI [0.16, 0.38]. In the present research, the subjective choices that researchers make in stimulus design have a substantial impact on observed effect sizes, but if a research team produces a large effect size for one research question, it does not necessarily mean that they will produce a large

effect size for another question. This pattern fails to support the hypothesis that some researchers have a flair for obtaining large and statistically significant results (see, e.g., Baumeister, 2016). Still, more research is needed on this point, since other research topics (e.g., stereotype threat, motivated reasoning), or more finely parsed subtopics, may yet yield evidence for expertise effects in conducting conceptual replications.⁶

⁶ We also reran these analyses, restricting the data to Hypotheses 3, 4, and 5, which are clearly within the same general area of research, moral psychology, to see whether we could find support for the flair hypothesis within a particular area of research. Once again, however, we found no evidence that observed effect sizes are predicted by the identity of the researchers that designed the materials (see Supplement 9 in the online supplemental materials for details).

Figure 1. (continued)

As might be expected, independent ratings of the quality of each study design (assessed in the Forecasting Study) were positively correlated with the obtained results. Higher quality sets of materials yielded larger observed effect sizes in the direction predicted by each original hypothesis (Cohen's d s), $r(62) = .31$, $p = .012$. Thus, it is possible that the inclusion of low-quality materials biases our analyses against finding support for hypotheses that are in fact true, when properly tested. We therefore repeated all of the meta-analytic analyses above, excluding 18 sets of materials that were rated as below 5 on a scale of 0 (*not at all informative*) to 10 (*extremely informative*) by independent raters in the Forecasting Study. As described in greater detail in Supplement 6 in the online supplemental materials, the results were substantively quite similar for all five hypotheses.

It is also possible that rather than artificially reducing the degree of observed support for a given hypothesis, lower quality materials

introduce psychometric artifacts such as poor reliability and validity which bias effects toward zero. We therefore further examined whether quality ratings predict larger effect size estimates in absolute terms, in other words larger estimates either consistent or inconsistent with the original hypothesis. Independent ratings of the quality of each study design were directionally positively correlated with the absolute value of the effect size estimates, but this relationship was not statistically significant, $r(62) = .20$, $p = .12$. Overall, the results suggest that the observed variability in effect sizes was not driven by a subset of lower quality study designs.

Aggregating results of the Main Studies and Replication Studies. Leveraging the combined samples of the Main Studies and Replication Studies allowed for more precise effect size estimates from each study version, as well as higher-powered estimates of the overall degree of support for each of the five original

Figure 1. (continued)

hypotheses. Aggregating all of the effect sizes across the two data collection efforts (Main Studies and Replication Studies) in random-effects meta-analyses ($k = 26$ for Hypotheses 1, 2, 3, and 5; $k = 24$ for Hypothesis 4) produced similar results to the separate meta-analyses above. Hypotheses 2 and 3 were supported ($d = 0.82$, 95% CI [0.55, 1.08], $p < .001$; $d = 0.29$, 95% CI [0.18, 0.39], $p < .001$). Hypothesis 5 was also associated with a statistically significant estimate in the expected direction, though, as above, the effect was negligible in size ($r = .04$, 95% CI [.01, .08], $p = .026$), leading to the conclusion that H5 was not empirically supported by the crowdsourced initiative. Later we report a Bayesian analysis casting further doubt on Hypothesis 5. Even under the null hypothesis significance testing framework, Hypotheses 1 and 4 were not supported ($d = 0.00$, 95% CI [-0.19, 0.19], $p = .997$, and $d = 0.05$, 95% CI [-0.02, 0.13], $p = .179$). We repeated these

analyses selecting only study versions rated as 5 or above in informativeness by the independent raters, (see Supplement 6 in the online supplemental materials), and nesting study (Main Studies vs. Replication Studies) within each hypothesis (see Supplement 9 in the online supplemental materials). Both of these additional analyses produced qualitatively similar results to the results above.

Comparing the results of the Main Studies and Replication Studies. As there is no single approach to determining whether an effect directly replicated or not (Brandt et al., 2014; Open Science Collaboration, 2015), we preregistered a number of criteria for whether the results of the Main Studies held up in the Replication Studies. These included correlating the Main Studies' and Replication Studies' effect sizes, comparing the statistical significance levels and direction of effects, and test-

Figure 1. (continued)

ing for statistically significant differences between the effect sizes from the Main Studies and the corresponding effect sizes from the Replication Studies. We further examined whether the effect was statistically significant after meta-analyzing across both the Main Studies and Replication Studies (see Figure 1), and we report a Bayesian analysis of differences in the Main Study and Replication results in Supplement 7 in the online supplemental materials.

Each of these criteria is an imperfect and incomplete measure of replication. For instance, a near perfect correlation in effect sizes could emerge even if replication effect sizes were dramatically smaller, so long as the rank ordering of effects remained consistent. Given such a pattern, it would be unreasonable to conclude the effects were robust and replicable. When it comes to comparing whether the replication effect is statistically significantly dif-

ferent from the original effect or not, this method is low in informational value when an original study has a statistically significant *p* value close to .05 with a lower bound of the confidence interval close to zero. With this *p* value, it is highly unlikely to find a statistically significant difference from the original result unless the replication point estimate is in the opposite direction of the original finding.

With these caveats in mind, we turn to comparing the results from the Main Studies and Replication Studies. In 51 of 64 cases (80%), the Replication Studies' effect was directionally consistent with the effect from the Main Studies. In 36 of those 51 cases (71%), when new participants were run using the same study design, statistically significant results were again statistically significant in the same direction, and nonsignificant effects were again nonsignificant. Further, 13 of 44 (30%) statistically signifi-

Table 2
Effect Sizes and Q , I^2 , and τ^2 Statistics From Meta-Analyses of Main Studies and Replication Studies

Hypothesis	Description	k	Effect size [95% CI]	Q	I^2 [95% CI]	τ^2 [95% CI]
Main studies						
1	Awareness of automatic prejudice	13	$d = .07 [-0.22, 0.37]$	$Q(12) = 897.51^{***}$	99.08% [98.20, 99.67]	0.28 [0.14, 0.81]
2	Extreme offers reduce trust	13	$d = 1.04 [0.61, 1.47]$	$Q(12) = 568.36^{***}$	98.25% [96.58, 99.36]	0.61 [0.31, 1.70]
3	Moral praise for needless work	13	$d = .33 [0.17, 0.50]$	$Q(12) = 152.45^{***}$	93.55% [87.39, 97.68]	0.09 [0.04, 0.26]
4	Proximal authorities drive legitimacy of performance enhancers	12	$d = .07 [-0.05, 0.20]$	$Q(11) = 89.72^{***}$	87.94% [75.82, 95.85]	0.04 [0.02, 0.13]
5	Deontological judgments predict happiness	13	$r = .06 [0.01, 0.11]$	$Q(12) = 52.91^{***}$	75.65% [52.68, 90.62]	0.01 [0.00, 0.02]
Replication studies						
1	Awareness of automatic prejudice	13	$d = -0.07 [-0.33, 0.19]$	$Q(12) = 773.19^{***}$	98.88% [97.81, 99.60]	0.23 [0.12, 0.64]
2	Extreme offers reduce trust	13	$d = .61 [0.32, 0.88]$	$Q(12) = 372.40^{***}$	97.09% [94.34, 98.98]	0.26 [0.13, 0.73]
3	Moral praise for needless work	13	$d = .24 [0.11, 0.38]$	$Q(12) = 129.49^{***}$	91.26% [82.81, 96.85]	0.05 [0.03, 0.16]
4	Proximal authorities drive legitimacy of performance enhancers	12	$d = .03 [-0.06, 0.12]$	$Q(11) = 47.45^{***}$	78.06% [55.84, 92.65]	0.02 [0.01, 0.07]
5	Deontological judgments predict happiness	13	$r = .03 [-0.04, 0.09]$	$Q(12) = 90.93^{***}$	86.39% [73.53, 94.97]	0.01 [0.00, 0.03]

*** $p < .001$.

cant findings from the Main Studies were not statistically significant in the Replication Studies. At the same time, six of 20 (30%) nonsignificant findings from the Main Studies were statistically significant in the Replication Studies.

We next examined whether effect sizes were significantly different in size between the two studies. We conducted z tests comparing each team-by-hypothesis combination across the two studies (e.g., Team 5's materials for Hypothesis 1 from the Main Studies, vs. Team 5's materials for Hypothesis 1 from the Replication Studies). Replication Studies' effect sizes were statistically significantly smaller than the corresponding effect in the Main Studies, according to z tests, in 21 of 64 cases, and statistically significantly larger in just one case, with no significant difference in 42 of 64 cases. This pattern agrees with the qualitative observation above that effect sizes tended to be somewhat smaller in the Replication Studies than in the Main Studies. This was quite unexpected—if anything, we anticipated that Mechanical Turk, as the less expensive, more expedient data source, might potentially yield smaller effect sizes. We can only speculate that the general decline effect across the two samples resulted from the slightly different populations of online respondents that were sampled, but the precise difference between the two samples that drove this result is unclear.

When directly replicated, a substantial minority of individual effect sizes reversed direction, changed significance levels across the $p < .05$ threshold, or were statistically significantly different from the initial result. At the same time, correlating the 64 effect sizes obtained in the Main Studies with the 64 effect sizes from the Replication Studies revealed very high correspondence between them in the aggregate, $r(62) = .92$, 95% CI [.88, .95], $p < .001$ (see Figure 2). Moreover, descriptively, the major overall findings from our Main Studies emerged in the Replication Studies as well. Effect sizes were again radically dispersed, with statistically significant effects in opposing directions obtained from different sets of materials designed to test three of the five research questions. Meta-analyzing across study versions, Hypotheses 2 and 3 were again supported, and Hypotheses 1 and 4 were not. The directional and statistically significant but very small estimate for Hypothesis

5 in the Main Studies was not statistically significant in the Replication Studies, yet also not meaningfully different in size (Gelman & Stern, 2006). Variability in effect sizes was again far more attributable to whether the hypothesis itself enjoyed overall support than to the skill of particular research teams at designing studies that returned large effects (see Supplement 9 in the online supplemental materials).

Publication bias analyses. We present funnel plots and the results of Egger's test (Egger, Davey Smith, Schneider, & Minder, 1997) for all of our meta-analytic results in Supplement 9 in the online supplemental materials. Because *all* of the study designs are reported in this article, there is, by definition, no publication bias in the results we have reported. Yet, we *did* find evidence of funnel plot asymmetries for Hypotheses 1, 2, and 5. As we discuss in greater detail in Supplement 9 in the online supplemental materials, these must reflect sample size effects that are idiosyncratic to the designs tested in this research. This result highlights one further advantage of crowdsourcing in comparison to the traditional practice of science: In a traditional meta-analysis of multiple studies conducted at different times, one cannot be certain whether funnel plot asymmetries reflect publication bias or some other sample size effect (see, e.g., Deeks, Macaskill, & Irwig, 2005), whereas in a crowdsourced project like this one, there is, by the very nature of the design, no publication bias.

Bayesian perspective on the results. Supplement 7 in the online supplemental materials provides an extended report of Bayesian analyses of the overall project results (the preregistered analysis plan is available at <https://osf.io/9jzy4/>). To summarize briefly, the Bayesian analyses find compelling evidence in favor of Hypotheses 2 and 3, moderate evidence against Hypothesis 1 and 4, and strong evidence against Hypothesis 5. Overall, two of five original hypotheses were confirmed aggregating across the different study designs. This pattern is generally consistent with the frequentist analyses reported above, with the exception that the frequentist approach suggests a very small but statistically significant ($p < .05$) effect in the direction predicted by Hypothesis 5 after aggregating across the different study designs, whereas the

Figure 2. Scatter plot comparing Main Study and Replication effect sizes (Cohen's d s). Each point in the scatter plot consists of one of 64 study designs. The continuous segment represents the fitted line; the dashed segment represents the 45-degree line. $H1$: Awareness of automatic prejudice, $H2$: Extreme offers reduce trust, $H3$: Moral praise for needless work, $H4$: Proximal authorities drive legitimacy of performance enhancers, $H5$: Deontological judgments predict happiness. See the online article for the color version of this figure.

Bayesian analyses find strong evidence *against* this prediction. The project coordinators, original authors who initially proposed Hypothesis 5, as well as further authors on this article concur with the Bayesian analyses that the effect is not empirically supported by the crowdsourcing hypotheses tests project, owing to the small estimate of the effect, and heterogeneity across designs. Regarding the main metascientific focus of this initiative, namely variability in results due to researcher choices, for all five hypotheses strong evidence of heterogeneity across different study designs emerged in the Bayesian analyses.

Forecasting Survey

We set up the forecasting survey to test whether scientists' predictions about the effect sizes and statistical significance levels (whether $p < .05$ or not) associated with the different sets of study materials would be positively correlated with the realized outcomes. Note that in asking forecasters to predict statistical significance levels, we are not endorsing the idea that something magical happens at $p = .05$, or the binary assumption of there being a result if $p < .05$ and none if $p > .05$ (Greenland, 2017). Yet, given that

in many fields and journals this criterion is used to indicate the minimum support required to claim an effect (see McShane & Gelman, 2017), we find that it is interesting to see whether a crowd of researchers can predict this binary outcome.

In addition, we tested whether monetary incentives or individual characteristics of the forecasters increased the accuracy of the predictions. The planned analyses for the forecasting study are detailed at <https://osf.io/9jzy4/>. Standard errors are clustered at two non-nested levels in all the regressions employing individual-level data: individual level and team-hypothesis version level (i.e., the level of a single study). Double clustering renders estimates robust to potential violations of independence among forecasts generated by the same individual over different versions of the study materials, and among predictions about the same set of study materials generated by different researchers.

Overall accuracy. To test our primary hypotheses regarding the accuracy of scientists' predictions, we examined whether there existed positive correlations between scientists' forecasts and the estimated effect sizes and statistical significance levels ($p < .05$ or not) from the different study versions in the Main Studies, at the

team-hypothesis version level. In addition, we performed paired t tests on aggregated prediction data and observed effect sizes to test whether scientists generally underestimated or overestimated the strength of each finding. As hypothesized, we observed a positive correlation between scientists' forecasts and the results being statistically significant in the predicted direction, $r(62) = 0.59$, 95% CI [0.40, 0.73], $p < .001$. The correlation between scientists' predictions and the observed effect sizes was likewise statistically significant: $r(62) = 0.71$, 95% CI [0.56, 0.81], $p < .001$.

We tested whether scientists underestimated or overestimated the realized outcomes by employing paired t tests between the vector collecting the average forecasts and the vectors collecting the effect sizes and directional statistical significance of each study version. Descriptively, for both effect sizes and directional statistical significance, predictions and outcomes were fairly aligned, with no differences reaching statistical significance. For directional statistical significance in terms of $p < .05$, the mean of the observed outcomes is $M = 0.58$ ($SD = 0.50$) and the mean of the forecasted outcomes is $M = 0.48$ ($SD = 0.09$), $t(63) = -1.78$, 95% CI of the difference of the means [-0.21, 0.01], $p = .080$. For effect sizes, the mean of the observed outcomes is $M = 0.31$ ($SD = 0.56$) and the mean of the forecasted outcomes is $M = 0.25$ ($SD = 0.10$), $t(63) = -1.02$, 95% CI of the difference of the means [-0.19, 0.06], $p = .311$. Evidence from the analysis of the forecasting survey supports the hypothesis that scientists' predictions are positively correlated with the realized outcomes, both in terms of effect sizes and in terms of whether the result is statistically significant or not for the different sets of study materials. Moreover, the analysis shows no evidence of systematic underestimation or overestimation of the realized outcomes.

Sensitivity to design choices. To test whether forecasters were sensitive to how different versions of the materials designed to test the same hypotheses affect research outcomes, we ran individual level regressions. These analyses tested whether scientists could predict results *within* each hypothesis, rather than only across them. The outcome (realized statistical significance in terms of $p < .05$, observed effect size) was the dependent variable and the individual prediction was the independent variable. As for all other individual level regressions, the standard errors were clustered at two non-nested levels: individual level (to account for the fact that each individual made several forecasts) and team-hypothesis version level (to account for the fact that the forecasts about the same set of materials might possibly be correlated). The model was estimated with either hypothesis fixed effects (exploiting only the variation in predictions across teams, as shown in equations [1a] and [1b]) or team fixed effects (exploiting only the variation in predictions across hypotheses, as shown in [2a] and [2b]).

$$SS_{ith} = \beta_0 + \beta_1 \hat{x}_{ith} + Hyp_h + \epsilon_{ith} \quad (1a)$$

$$EE_{ith} = \beta_0 + \beta_1 \hat{x}_{ith} + Hyp_h + \epsilon_{ith} \quad (1b)$$

The dependent variables SS_{ith} and EE_{ith} are the realized outcomes, the dummy variable being positive if the study is statistically significant in (1a), and realized effect size in (1b), respectively. The independent variables are the individuals' forecasts, x_{ith} for the predictions regarding statistical significance in terms of $p < .05$ and $\hat{x}_{ith}$ for the predictions regarding effect size. Hyp_h identify the hypothesis fixed effects, and $Team_t$ are the team fixed effects.

$$SS_{ith} = \beta_0 + \beta_1 x_{ith} + Team_t + \epsilon_{ith} \quad (2a)$$

$$EE_{ith} = \beta_0 + \beta_1 \hat{x}_{ith} + Team_t + \epsilon_{ith} \quad (2b)$$

Separately including only hypothesis or only team fixed effects allows us to test whether the forecasts are associated with the realized outcomes using only the variation in forecasts within hypotheses (making predictions for the different teams within hypotheses) or only the variation in forecasts within teams (making predictions for the different hypotheses within teams).

The individual prediction coefficient was statistically significant in the expected direction in both the regressions with only hypothesis fixed effects, and in the regressions with only team fixed effects. This holds for predicting both statistical significance levels ($\beta_1 = .148$, $t[9018] = 4.07$, $p < .001$, controlling for hypotheses, $\beta_1 = .255$, $t[9007] = 4.38$, $p < .001$ controlling for teams), and effect sizes ($\beta_1 = 0.097$, $t[9018] = 2.16$, $p = .031$, controlling for hypotheses, $\beta_1 = 0.228$, $t[9007] = 2.68$, $p = .007$ controlling for teams), and shows that forecasters were able to anticipate results from different teams of materials designers within each hypothesis, as well as different hypotheses within each team of materials designers. For completeness, we also estimated the results without any fixed effects ($\beta_1 = 0.309$, $t[9022] = 43.04$ for predictions on whether the result is statistically significant [$p < .05$] or not, $\beta_1 = 0.309$, $t[9022] = 2.38$ for predictions regarding effect size) and with both team and hypotheses fixed effects ($\beta_1 = 0.089$, $t[9003] = 2.78$ for predictions whether the result is statistically significant [$p < .05$] or not, $\beta_1 = 0.091$, $t[9003] = 2.77$ for predictions regarding effect size), and the individual prediction coefficient is statistically significant in these models as well (see Tables S5.4a and S5.5 in Supplement 5 in the online supplemental materials).⁷ Furthermore, we estimated equations (1a) and (2a) as a probit model (see Table S5.4b in Supplement 5 in the online supplemental materials), obtaining similar results as those obtained using the linear probability model. In short, scientists were able to predict not only which hypotheses would receive empirical support (see Figure 3a) but also variability in results for the same hypothesis based on the design choices made by different research teams (see Figure 3b).

We report several further analyses of the Forecasting Study in Supplement 5 in the online supplemental materials, for the interested reader. In particular, we examine whether monetary incentives increase the accuracy of forecasts (they do not, at least with the relatively small incentives on offer in this study), whether characteristics of the forecaster, such as job rank and confidence in their forecasts, predict accuracy (they do not consistently do so), and repeat our primary analyses for the data from the Replication Studies and aggregating across the Main Studies and Replication Studies (the results are similar to those reported here).

⁷ In all four models, there was a statistically significant association between individual forecasts and outcomes. This is true for both the predictions regarding whether the study will find a statistically significant effect in the hypothesized direction and the predictions regarding the realized effect size. Note however that, because the independent variable (i.e., the individual forecasts) is likely to be measured with error, the estimated coefficients reported in this paragraph are potentially biased downwards. Measurement error would artificially reduce the correspondence between forecasts and outcomes, leading to a conservative test of forecaster accuracy.

Figure 3. (a) Correlation between average predicted effect size and observed effect size for each study design. The continuous segment represents the fitted line; the dashed segment represents $y = x$. *H1*: Awareness of automatic prejudice, *H2*: Extreme offers reduce trust, *H3*: Moral praise for needless work, *H4*: Proximal authorities drive legitimacy of performance enhancers, *H5*: Deontological judgments predict happiness. (b) Correlation between average predicted effect size and observed effect size for each version of the study materials, separately for each of the five hypotheses. Continuous segments represent fitted lines; dashed segments represent $y = x$. *H1*: Awareness of automatic prejudice, *H2*: Extreme offers reduce trust, *H3*: Moral praise for needless work, *H4*: Proximal authorities drive legitimacy of performance enhancers, *H5*: Deontological judgments predict happiness. See the online article for the color version of this figure.

Discussion

How contingent is support for scientific hypotheses on the subjective choices that researchers make when designing studies? Concerns about the potential dependency of findings on the stimuli used to capture them have been raised repeatedly (e.g., Baribault et al., 2018; Campbell & Fiske, 1959; Judd et al., 2012; Monin & Oppenheimer, 2014; Monin et al., 2007; Wells & Windschitl, 1999). In contrast, the extent to which this problem presents a challenge to conducting research investigations and interpreting research findings has never been directly examined. In this crowdsourced project, when up to 13 independent research teams designed their own studies to test five original research questions, variability in observed effect sizes proved dramatic, with the Bayesian analyses confirming overwhelming evidence of heterogeneity for four of five hypotheses and compelling evidence in the fifth case (see Supplement 7 in the online supplemental materials). Descriptively, different research teams designed studies that returned statistically significant effects in opposing directions for the same research question for four of five hypotheses in the Main Studies, and three of five hypotheses in the Replication Studies (see Supplement 9 in the online supplemental materials). In other words, even when some or most teams created studies that substantiated a theoretical prediction, at least one other team's design found the opposite. Even the most consistently supported original hypotheses still exhibited a wide range of effect sizes, with the smallest range being $d = -0.37$ to $d = 0.26$ (Hypothesis 4, Replication Studies). Although the hypothesis being tested explained substantial variability in effect sizes (i.e., some hypotheses received more consistent support than others), there remained substantial unexplained heterogeneity after accounting for the hypothesis being tested, implying that idiosyncratic choices in stimulus design have a very large effect on observed results, over and above the overall support (or lack thereof) for the hypothesis in question.

Crowdsourcing makes more transparent the true consistency of support for a scientific prediction, and provides the opportunity to leverage the collective experience and perspectives of a crowd of scientists via aggregation (Bates & Granger, 1969; Galton, 1907; Lorenz, Rauhut, Schweitzer, & Helbing, 2011; Silberzahn et al., 2018; Surowiecki, 2004). Meta-analytically combining effect sizes across the various conceptual replications yielded overall support for two of five of the original predictions, and a Bayesian analysis likewise supported two of five hypotheses. Crowdsourcing hypothesis tests can confirm and disconfirm predictions in a convincing way, by providing converging evidence across independent investigators who are unbiased by each other's approaches or knowledge of the original finding.

Contrary to the *flair hypothesis* (Baumeister, 2016) that some researchers are more adept at obtaining empirical support for their predictions, none of the 15 different teams involved in this project designed studies associated with more consistent support for the original ideas. This noneffect occurred despite variable seniority of team leaders, who ranged from doctoral students to chaired full professors, with citation counts ranging from zero into the tens of thousands. The present findings further suggest that replication results are more attributable to the robustness and generalizability of the original finding than the skill of the scientist carrying out the replication (whether a direct or conceptual replication; Bench, Rivera, Schlegel, Hicks, & Lench, 2017; Open Science Collaboration, 2015). Although replicating some studies certainly requires specialized technical knowledge (e.g., of neuroimaging technology), evidence that

disappointing reproducibility rates for published research (e.g., Dewald, Thursby, & Anderson, 1986; Klein et al., 2014; LeBel, 2015; Open Science Collaboration, 2015) are attributable to a dearth of replicator competence remains lacking. That said, further metascientific work is needed on the role of expertise in replication results (Tierney, Ebersole, Hardy, Gordon, et al., 2019).

A substantial degree of variability in the results was accounted for by the original hypotheses themselves, which—as noted earlier—differed in their overall empirical support (see Figure 1). Although the original effects all replicated using the original materials (when combining the results of the Main Studies and Replication Studies), three effects were unsupported overall in the alternative study designs, in some cases returning estimates in the opposite direction than predicted.

As confirmed in the Bayesian analyses of the project results (Supplement 7 in the online supplemental materials), all five original hypotheses exhibited wide variability in support across different study designs. Although the present project was able to parse the two, in typical research contexts this heterogeneity in results attributable to study design choices coexists and potentially interacts with heterogeneity in results because of population differences (McShane et al., 2019; Tiokhin et al., 2019). Discrepant results and variability in research findings (Open Science Collaboration, 2015; Schweinsberg et al., 2016) are perhaps unavoidable, and might best be embraced as a normal aspect of the scientific process. In terms of building solid theory, it may be necessary to vary stimuli and study designs (Baribault et al., 2018; Caruso et al., 2017; the present initiative), employ a variety of statistical specifications (Silberzahn et al., in 2018; Simonsohn, Simmons, & Nelson, 2016; Steegen, Tuerlinckx, Gelman, & Vanpaemel, 2016), and replicate findings across more geographic locations and populations (Henrich, Heine, & Norenzayan, 2010), before drawing definitive conclusions. With regard to communicating findings both within and outside the scientific community, more conservative messaging regarding new research conducted in a single population or relying heavily on a specific experimental paradigm seems warranted.

Implications for the Five Original Hypotheses

The primary goal of this initiative was to examine effect size dispersion when independent investigators design studies to address the same research questions. A secondary purpose was to evaluate the evidence for the five original effects targeted in the crowdsourced conceptual replications. Below we assess current support and potential future directions for Hypothesis 1–5, in consultation with the original team that volunteered each research idea for the initiative.

Hypothesis 1: Awareness of automatic prejudice. This effect directly replicated using the original Uhlmann and Cunningham (2000) questionnaire items, with participants in both the Main Study and Replication expressing overall agreement to the items “Although I don’t necessarily agree with them, I sometimes have prejudiced feelings (like gut reactions or spontaneous thoughts) that I don’t feel I can prevent,” and “At times stereotypical thoughts about minorities coming into my head without my necessarily intending them to.” As in the original data collections by Uhlmann and Cunningham (2000), mean responses to these items were significantly above the neutral scale midpoint of four (1 = *strongly disagree*, 4 = *neutral*, 7 = *strongly agree*). At the same

time, conceptual replications by different research teams employing alternative questions failed to confirm the hypothesis that participants so openly self-report automatic prejudices. Aggregating across the different study designs via meta-analysis reveals no statistically significant effect in the expected direction, and a Bayesian analysis found moderate evidence against H1. On reflection, the double-barreled nature of the original items, invoking both lack of intentions and prejudiced reactions, as well as the use of qualifiers (“I sometimes,” “At times”) might have biased participants’ responses toward agreement. Further shortcomings of the original study design are the lack of a relative comparison group (e.g., nonminorities and members of dominant groups such as White men), and the absence of any probe items regarding positive or favorable thoughts.

In sum, the present initiative to crowdsource hypothesis tests casts serious doubt on whether overall endorsement of self-perceived automatic prejudice is generally as high as initially reported by Uhlmann and Cunningham (2000). Yet, it does not call into question evidence that different measures of beliefs are correlated at an individual level with scores on implicit measures of attitude (Hahn et al., 2014) and that awareness of automatic associations can be experimentally increased (Hahn & Gawronski, 2019). As of yet there are no systematic reviews or meta-analyses on the empirical relationships between awareness indices and automatic associations. From the present crowdsourced project, we cannot conclude that everyday people believe themselves to be as biased as implicit and indirect measures of automatic associations suggest they are. Indeed, the present results, relying on a wide array of study designs, suggest they do not generally see themselves as implicitly prejudiced. Opportunities to improve validated self-report measures of beliefs about one’s automatic prejudices toward various social groups, and to use them as predictors and outcome measures in future investigations, remain open.

Hypothesis 2: Extreme offers reduce trust. This crowdsourcing initiative found consistent evidence for Hypothesis 2 across the range of conceptual replications, as well as in the direct replications using the original materials. Both frequentist and Bayesian analyses supported this particular prediction, with the Bayesian analyses confirming compelling evidence for this hypothesis despite heterogeneity in estimates across different study designs. This result is consistent with a recent meta-analysis (Hüffmeier, Freund, Zerres, Backhaus, & Hertel, 2014), which found that hardline negotiation tactics (of which extreme first offers are one example) are associated with more negative socioemotional outcomes in negotiations (i.e., perceptions that the hardline negotiator is unreasonable and uncooperative). However, this meta-analysis did not *specifically* examine extreme first offers or trust. Although our findings provide initial support for the idea that extreme first offers indeed reduce trust on the part of the recipient, that this reduced trust consequently diminishes information exchange and value creation remains to be demonstrated. It also remains unclear to what extent such effects generalize across cultures. Given that negotiators in some cultural settings may be more accustomed to receiving extreme first offers than negotiators in other cultural settings, this effect may indeed be culturally moderated. This possibility is currently being examined in an ongoing international replication project (Schweinsberg et al., 2019) that will assess the cultural boundary conditions of this effect.

Hypothesis 3: Moral praise for needless work. Earlier findings that Americans morally praise individuals who continue at their

job after coming into sudden wealth were likewise confirmed by the crowdsourced initiative. Aggregating via meta-analysis across distinct studies independently created by different research teams, both the frequentist and Bayesian analyses find compelling evidence in favor of the needless work hypothesis. Although the robustness of the effect to different operationalizations is now confirmed in two large U.S. samples via the present host of conceptual replications, the original hypothesis of cross-cultural variability has yet to be put to a rigorous empirical test. The original research predicted that praise for those who work in the absence of any material need is steeped in the Protestant work ethic, and hence should be strongest among those with greater degrees of exposure to U.S. culture (Poehlman, 2007; Uhlmann et al., 2009).

Because there is no systematic literature review or meta-analysis on this topic, an ongoing crowdsourced project by Tierney, Ebersole, Hardy, Chapman, et al. (2019) will attempt to directly replicate this and other original findings regarding work morality across four countries (the United States, the United Kingdom, Australia, and India). Relying on a creative destruction approach to replication, the Tierney, Ebersole, Hardy, Chapman, et al. (2019) initiative will pit the original prediction that moral praise for needless work only characterizes U.S. culture against theories positing the general moralization of work across cultures, regional differences within the United States (i.e., New England vs. other regions; Fisher, 1989), and valorization of work as a means of personal fulfillment in postmaterialist societies (Inglehart, 1997; Inglehart & Welzel, 2005). Thus, further facets of the robustness, generalizability, and potential cultural boundedness of this effect remain to be explored in future research. For now, we conclude that aggregating across the crowdsourced study designs, the needless work hypothesis is supported for U.S. participants, but the originally hypothesized moderation by culture (Poehlman, 2007; Uhlmann et al., 2009) remains to be demonstrated.

Hypothesis 4: Proximal authorities drive legitimacy of performance enhancers. The original finding that the dictates of proximal authorities (e.g., the league, the competitive circuit) have a larger impact on judgments of the acceptability of using performance enhancing drugs (PEDs) than the law was not supported in this crowdsourced initiative. Although the finding directly replicated using the original materials, across a dozen different, independently developed study designs, people were not more opposed to the use of PEDs when they are banned by a proximal authority than when they are illegal, and the Bayesian analysis found moderate evidence against this hypothesis. This result concurs with follow-up studies done by the research team who contributed this hypothesis (Landy et al., 2017), which were conducted after this project began. These subsequent studies find that both types of authority contribute to normative judgments of PED use, to similar degrees. There is currently no systematic review or meta-analysis of judgments of PED use, but Landy et al. (2017) employed an exploratory, deep-dive methodology, in which they tested 11 different potential explanations for opposition to the use of these substances. They concluded that PED use is opposed for three primary reasons: it violates moral norms of fairness, it poses a risk of harm to the user, and it tends to violate legitimate conventional rules. The present results help to clarify this last reason, by showing that the precise source of those rules—the law or a more proximal authority—does not affect levels of opposition.

Hypothesis 5: Deontological judgments predict happiness.

Although the original pattern of results once again directly replicated using the original materials, the hypothesis that individuals who tend to make deontological (vs. utilitarian) judgments report different levels of personal happiness was not supported overall by the crowdsourced conceptual replications. Although a statistically significant directional effect in support of H5 was reported in the Main Studies, the aggregated estimate was close to zero, and the effect did not reach statistical significance in the Replication Studies. Overall, the Bayesian analysis found strong evidence *against* this original prediction. There has not previously been a systematic review or meta-analysis of the relationship between moral stance and happiness, though prior research has linked both processes to emotional and intuitive responding (e.g., Everett et al., 2016; Greene, 2013; Lieberman, 2013; Phillips et al., 2017; Singer, 2005). These results fail to find support for an association between deontological moral judgments and hedonic happiness that has been suggested—although not empirically confirmed—by this prior work. Although laypeople appear to believe that part of what brings happiness is living a moral life (Phillips et al., 2017), adherence to deontological versus utilitarian ethical principles does not seem to relate to one's overall happiness.

Forecasting Findings

Scientists can predict whether a published finding will replicate from the research reports (Camerer et al., 2016; Dreber et al., 2015) and benchmark findings plus the materials for further experimental conditions (DellaVigna & Pope, 2018a, 2018b). We find that examination of the materials for an unpublished study is sufficient for scientists to successfully anticipate the outcome. In our forecasting survey, predictions by independent scientists were significantly correlated with both effect sizes and whether the observed results were statistically significant in the hypothesized direction, and the average predictions were similar to the observed outcomes. Monetary incentives failed to improve forecasters' predictive performance. Although speculative, it is possible that scientists who opted into and completed an extensive survey about predicting research findings were sufficiently intrinsically motivated to be accurate, so external incentives did not further increase their motivation (see Lakhani & Wolf, 2005, and Lakhani, Jeppesen, Lohse, & Panetta, 2007, regarding the tendency for crowdsourced initiatives to leverage intrinsic motivations). Another potential explanation is that the financial incentives (up to \$60) were not sufficiently strong to affect accuracy.

Comparatively more senior academics (in terms of job rank) were more accurate at forecasting statistical significance levels (i.e., whether the study's outcome would be $p < .05$ in the predicted direction or not), but not effect sizes (see Supplement 5 in the online supplemental materials). Other indices of scientific eminence, such as number of peer reviewed publications, were unrelated to forecasting accuracy. In a separate investigation, DellaVigna and Pope (2018a) found that more senior academics (in terms of job rank and citations), if anything, underperformed junior academics at predicting how different incentives would influence the effort and performance of experimental subjects. Moreover, academics in general did no better than lay people (undergraduates, MBA students, and MTurk workers) at rank-ordering the effectiveness of different experimental treatments (DellaVigna & Pope, 2018a). More research is needed on whether

traditional indices of scientific eminence (Sternberg, 2016; Vazire, 2017) are associated with any advantage in designing or predicting the results of scientific studies.

Unique to the present study, we show that independent scientists are not only able to predict study results with some success by merely examining the materials, but are also sensitive to how design choices influence the degree of empirical support for a specific claim. Forecasters predicted research results with significant accuracy not just across but also within each of the five hypotheses. This suggests some fine-grained sensitivity to how different operationalizations of the same hypothesis can impact results. More forecasting surveys and other tools aggregating beliefs such as prediction markets are needed to determine the accuracy of scientists' intuitions about how contextual factors affect research outcomes—for instance, whether scientists are able to anticipate cultural differences in effects, and whether specializing in research on culture confers any special advantage. Ongoing projects from our group examine whether academics can predict the heterogeneity statistics in replication results for prime-to-behavior effects (Tierney, Ebersole, Hardy, Gordon, et al., 2019), differences in replication effect sizes when the same experiment is run in multiple laboratories (Schweinsberg et al., 2019), and whether findings from the field of strategic management generalize to other time periods and geographies (Delios et al., 2019).

Limitations and Future Directions

This project represents an early foray into the crowdsourcing of stimulus selection and study designs (see also Baribault et al., 2018), with important limitations that should be addressed in future initiatives. The primary metascientific purpose of this initiative was to examine the impact of scientists' design choices on effect size estimates. Still, a number of aspects of our approach may have led to artificial homogeneity in study designs. In particular, materials designers were restricted to creating simple experiments with a self-reported dependent measure that could be run online in five minutes or less. Further, the key statistical test of the hypothesis had to be a simple comparison between two conditions (for Hypotheses 1–4), or a Pearson correlation (for Hypothesis 5). Full 30-min- to hour-long-laboratory paradigms with factorial designs, research confederates, and more complex manipulations and outcome measures (e.g., behavioral measures) contain far more researcher choice points and may be associated with even greater heterogeneity in research results. In addition, the project coordinators recruited the materials designers from their own social networks, potentially biasing the project toward demographic and intellectual homogeneity (Ibarra, 1995, 1997). Future initiatives should recruit materials designers more broadly, to better represent the diversity of perspectives within a field or subfield (Duarte et al., 2015; McGuire, 1973; Monin et al., 2007).

Another limitation is that our participants all participated in tests of multiple research questions. In meta-analysis, it is typically assumed that all samples are independent of one another, but this assumption is violated in our data. This assumption is not problematic in the univariate meta-analyses that we present in the main text, but it does complicate the multivariate meta-analysis we present in Supplement 8 in the online supplemental materials. Future crowdsourced initiatives could

perhaps assign each participant to only *one* research design, or focus exclusively on a single research question, to avoid these participant-level correlations across hypotheses, which are not accounted for in our primary analyses. This would allow for a straightforward multivariate meta-analytic approach, in which participants are nested within designs, which are nested within hypotheses. Each research team was also free to develop their own dependent measures, which meant that we could not directly compare raw results across different designs, but could only compare standardized effect sizes. Future projects in this vein might constrain dependent measures to allow clean, straightforward comparisons of the effects of multiple, independently developed experimental manipulations.

This initiative to crowdsource hypothesis tests also targeted only five original hypotheses, leaving us unable to identify which features of a research idea might be associated with more or less heterogeneity in study designs and outcomes. Some research ideas may naturally feature a greater latitude of construal (Beck, McCauley, Segal, & Hershey, 1988; Dunning, Meyerowitz, & Holzberg, 1989), leading different teams to create more varied experimental paradigms in order to test them. In the extreme, hypotheses that are theoretically underspecified (unlike the present H1–H5) may result in a chaos of operationalizations as the materials designers impose their own priors and assumptions on the idea. Thus, one way to reduce the role of subjective researcher choices in research outcomes may be to more fully flesh out the underlying theory at the outset (Dijksterhuis, 2014; McGuire, 1973; Stroebe & Strack, 2014).

Our limited number of target hypotheses also means that one cannot generalize the present results to all hypotheses in all subfields. We cannot conclude that only 40% of research ideas that directly replicate will be supported in conceptual replications, or that for the majority of research questions different designs will return statistically significant effects in opposing directions. Those are the results of this project only, and further initiatives to crowdsource hypothesis tests are needed before drawing definitive conclusions about the impact of subjective researcher choices on empirical outcomes.

Perhaps the most concrete methodological limitation of the present project was the modest sample of forecasters ($N = 141$), which reduced the statistical power of the relevant analyses. Our sample size was comparable with those for prior surveys examining the forecasting abilities of academics. For example, DellaVigna and Pope (2018a, 2018b) recruited 208 academics for their forecasting research, Dreber et al. (2015) had 47 and 45 active traders in their two prediction markets for replications, Camerer et al.'s (2016) prediction market had 97 participants, Forsell et al. (2019) included 78 participants, and Camerer et al. (2018) featured two conditions with 114 and 92 participants in each treatment. For the present project, we recruited the largest sample we could, given our forecasters' massive task of reviewing, making quality assessments, and predicting the results from 64 distinct sets of experimental materials. Still, the relatively small group of forecasters in our survey indeed limits our conclusions. Furthermore, there may be overlap between the samples of forecasters included in this and other studies, as they were recruited by similar methods. Further research is needed using higher-powered designs, especially with regard to the potential role of forecaster characteristics in moderating predictive accuracy.

Finally, the crowdsourcing hypothesis tests approach shares certain costs and benefits with other crowd approaches to scientific research (Uhlmann et al., 2019). In comparison with the standard

approach of relying on a small team, recruiting a crowd of collaborators enables big science, democratizes access to projects, and more effectively assesses the robustness of the findings. Yet at the same time, crowdsourcing study designs are inefficient, in that for the same effort and expense, initial evidence for a far greater number of interesting ideas could have been obtained using a small team or solo investigator approach. In future work, the return on investment from crowdsourcing hypothesis tests may be greatest for theoretically important findings that are well established with a specific paradigm, and whose robustness to alternative methodological approaches is of general interest.

Conclusions

The present crowdsourced project illustrates the dramatic consequences of researcher design choices for scientific results. This initiative also provides a roadmap for future crowdsourced approaches to testing the generality of scientific theories. If a scientific prediction is theoretically important enough, or has practically significant policy and societal implications, future investigations could assign it to multiple laboratories to independently operationalize and carry out empirical tests. The extent to which the results converge (and diverge) across investigations can then be used to inform discussion and debate, revise theory, and formulate policy.

Scientists craft theories with the ambitious goal of unifying potentially disparate findings into coherent, generalizable structures of knowledge. This process is often arduous and lengthy and may be impeded by features of the standard approach to scientific inquiry. Nonetheless, this process can be streamlined through collective action. As the present investigation demonstrates, bringing many perspectives and operationalizations to bear on hypotheses provides a richer account of phenomena than would occur if researchers and teams worked in isolation. Moreover, we also showed that independent researchers are able to identify not only the hypotheses that are more likely to be supported by an empirical investigation but also the research designs that, within a specific hypothesis, are more likely to lead to significant effect. This suggests that researchers can determine the features of the hypotheses, of the methods, and of the research designs that are systematically associated with the effect size and the statistical significance of a research question. Through crowdsourced collaborations such as this one, researchers can craft theories with more confidence and better understand just how far they extend.

References

- Alogna, V. K., Attaya, M. K., Aucoin, P., Bahník, Š., Birch, S., Birt, A. R., . . . Zwaan, R. A. (2014). Registered replication report: Schooler & Engstler-Schooler (1990). *Perspectives on Psychological Science*, 9, 556–578. <http://dx.doi.org/10.1177/1745691614545653>
- Ames, D. R., & Mason, M. F. (2015). Tandem anchoring: Informational and politeness effects of range offers in social exchange. *Journal of Personality and Social Psychology*, 108, 254–274. <http://dx.doi.org/10.1037/pspi0000016>
- Annas, J. (1993). *The morality of happiness*. New York, NY: Oxford University Press.
- Aristotle. (2002). *Nicomachean ethics* (S. Brodie & C. Rowe, Eds.). Oxford, UK: Oxford University Press. (Original work published 340 BCE)
- Aronson, E., & Mills, J. (1959). The effect of severity of initiation on liking for a group. *The Journal of Abnormal and Social Psychology*, 59, 177–181. <http://dx.doi.org/10.1037/h0047195>

- Banaji, M. R. (2001). Implicit attitudes can be measured. In H. L. Roediger, III, J. S. Nairne, I. Neath, & A. M. Surprenant (Eds.), *The nature of remembering: Essays in honor of Robert G. Crowder* (pp. 117–150). Washington, DC: American Psychological Association.
- Baribault, B., Donkin, C., Little, D. R., Trueblood, J. S., Oravecz, Z., van Ravenzwaaij, D., . . . Vandekerckhove, J. (2018). Metastudies for robust tests of theory. *Proceedings of the National Academy of Sciences of the United States of America*, 115, 2607–2612. <http://dx.doi.org/10.1073/pnas.1708285114>
- Bartko, J. J. (1976). On various intraclass correlation reliability coefficients. *Psychological Bulletin*, 83, 762–765. <http://dx.doi.org/10.1037/0033-2909.83.5.762>
- Bates, J. M., & Granger, C. W. J. (1969). The combination of forecasts. *Operational Research Quarterly*, 20, 451–468. <http://dx.doi.org/10.1057/jors.1969.103>
- Baumeister, R. F. (2016). Charting the future of social psychology on stormy seas: Winners, losers, and recommendations. *Journal of Experimental Social Psychology*, 66, 153–158. <http://dx.doi.org/10.1016/j.jesp.2016.02.003>
- Beck, L., McCauley, C., Segal, M., & Hershey, L. (1988). Individual differences in prototypicality judgments about trait categories. *Journal of Personality and Social Psychology*, 55, 286–292. <http://dx.doi.org/10.1037/0022-3514.55.2.286>
- Belsie, L. (2011). Powerball numbers: Why do lottery winners keep working? *The Christian Science Monitor*. Retrieved from <https://www.csmonitor.com/Business/new-economy/2011/0602/Powerball-numbers-Why-do-lottery-winners-keep-working>
- Bench, S. W., Rivera, G. N., Schlegel, R. J., Hicks, J. A., & Lench, H. C. (2017). Does expertise matter in replication? An examination of the Reproducibility Project: Psychology. *Journal of Experimental Social Psychology*, 68, 181–184. <http://dx.doi.org/10.1016/j.jesp.2016.07.003>
- Bentham, J. (1970). *An introduction to the principles of morals and legislation*. London, UK: Althone Press. (Original work published 1823)
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. Chichester, UK: Wiley. <http://dx.doi.org/10.1002/9780470743386>
- Borenstein, M., Hedges, L. V., Higgins, J. P., & Rothstein, H. R. (2010). A basic introduction to fixed-effect and random-effects models for meta-analysis. *Research Synthesis Methods*, 1, 97–111. <http://dx.doi.org/10.1002/jrsm.12>
- Brandt, M. J., Ijzerman, H., Dijksterhuis, A., Farach, F., Geller, J., Giner-Sorolla, R., . . . van 't Veer, A. (2014). The replication recipe: What makes for a convincing replication? *Journal of Experimental Social Psychology*, 50, 217–224. <http://dx.doi.org/10.1016/j.jesp.2013.10.005>
- Brehm, J. W. (1956). Postdecision changes in the desirability of alternatives. *The Journal of Abnormal and Social Psychology*, 52, 384–389. <http://dx.doi.org/10.1037/h0041006>
- Camerer, C. F., Dreber, A., Forsell, E., Ho, T. H., Huber, J., Johannesson, M., . . . Wu, H. (2016). Evaluating replicability of laboratory experiments in economics. *Science*, 351, 1433–1436. <http://dx.doi.org/10.1126/science.aaf0918>
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., . . . Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2, 637–644. <http://dx.doi.org/10.1038/s41562-018-0399-z>
- Campbell, D. T., & Fiske, D. W. (1959). Convergent and discriminant validation by the multitrait-multimethod matrix. *Psychological Bulletin*, 56, 81–105. <http://dx.doi.org/10.1037/h0046016>
- Caruso, E. M., Shapira, O., & Landy, J. F. (2017). Show me the money: A systematic exploration of manipulations, moderators, and mechanisms of priming effects. *Psychological Science*, 28, 1148–1159. <http://dx.doi.org/10.1177/0956797617706161>
- Chandler, J., Mueller, P., & Paolacci, G. (2014). Nonnaïveté among Amazon Mechanical Turk workers: Consequences and solutions for behavioral researchers. *Behavior Research Methods*, 46, 112–130. <http://dx.doi.org/10.3758/s13428-013-0365-7>
- Chandler, J., Paolacci, G., & Mueller, P. (2013). Risks and rewards of crowdsourcing marketplaces. In P. Michelucci (Ed.), *Handbook of human computation* (pp. 377–392). New York, NY: Sage. http://dx.doi.org/10.1007/978-1-4614-8806-4_30
- Coffman, L., & Niehaus, P. (2014). *Pathways of persuasion* [Working paper]. Retrieved from <https://econweb.ucsd.edu/~pniehaus/papers/persuasion.pdf>
- Crandall, C. S., & Sherman, J. W. (2016). On the scientific superiority of conceptual replications for scientific progress. *Journal of Experimental Social Psychology*, 66, 93–99. <http://dx.doi.org/10.1016/j.jesp.2015.10.002>
- Deeks, J. J., Macaskill, P., & Irwig, L. (2005). The performance of tests of publication bias and other sample size effects in systematic reviews of diagnostic test accuracy was assessed. *Journal of Clinical Epidemiology*, 58, 882–893. <http://dx.doi.org/10.1016/j.jclinepi.2005.01.016>
- Delios, A., Tan, H., Wu, T., Wang, Y., Viganola, D., Gordon, M., . . . Uhlmann, E. (2019). *Can you step into the same river twice? Examining the context sensitivity of research findings from archival data*. Manuscript submitted for publication.
- DellaVigna, S., & Pope, D. G. (2018a). Predicting experimental results: Who knows what? *Journal of Political Economy*, 126, 2410–2456. <http://dx.doi.org/10.1086/699976>
- DellaVigna, S., & Pope, D. (2018b). What motivates effort? Evidence and expert forecasts. *The Review of Economic Studies*, 85, 1029–1069. <http://dx.doi.org/10.1093/restud/rdx033>
- Dewald, W. G., Thursby, J. G., & Anderson, R. G. (1986). Replication in empirical economics: The journal of money, credit and banking project. *The American Economic Review*, 76, 587–603.
- Diener, E., Emmons, R. A., Larsen, R. J., & Griffin, S. (1985). The satisfaction with life scale. *Journal of Personality Assessment*, 49, 71–75. http://dx.doi.org/10.1207/s15327752jpa4901_13
- Dijksterhuis, A. (2014). Welcome back theory! *Perspectives on Psychological Science*, 9, 72–75. <http://dx.doi.org/10.1177/1745691613513472>
- Dodge, T., Williams, K. J., Marzell, M., & Turrissi, R. (2012). Judging cheaters: Is substance misuse viewed similarly in the athletic and academic domains? *Psychology of Addictive Behaviors*, 26, 678–682. <http://dx.doi.org/10.1037/a0027872>
- Dovidio, J. F., Kawakami, K., & Gaertner, S. L. (2002). Implicit and explicit prejudice and interracial interaction. *Journal of Personality and Social Psychology*, 82, 62–68. <http://dx.doi.org/10.1037/0022-3514.82.1.62>
- Doyen, S., Klein, O., Simons, D. J., & Cleeremans, A. (2014). On the other side of the mirror: Priming in cognitive and social psychology. *Social Cognition*, 32, 12–32. <http://dx.doi.org/10.1521/soco.2014.32.suppl.12>
- Dreber, A., Pfeiffer, T., Almenberg, J., Isaksson, S., Wilson, B., Chen, Y., . . . Johannesson, M. (2015). Using prediction markets to estimate the reproducibility of scientific research. *Proceedings of the National Academy of Sciences of the United States of America*, 112, 15343–15347. <http://dx.doi.org/10.1073/pnas.1516179112>
- Duarte, J. L., Crawford, J. T., Stern, C., Haidt, J., Jussim, L., & Tetlock, P. (2015). Political diversity will improve social and personality psychological science. *Behavioral and Brain Sciences*, 38, 1–13. <http://dx.doi.org/10.1017/S0140525X14000430>
- Dunaway, B., Edmonds, A., & Manley, D. (2013). The folk probably do think what you think they think. *Australasian Journal of Philosophy*, 91, 421–441. <http://dx.doi.org/10.1080/00048402.2012.724430>
- Dunning, D., Meyerowitz, J. A., & Holzberg, A. D. (1989). Ambiguity and self-evaluation: The role of idiosyncratic trait definitions in self-serving assessments of ability. *Journal of Personality and Social Psychology*, 57, 1082–1090. <http://dx.doi.org/10.1037/0022-3514.57.6.1082>
- Earp, B. D. (in press). Falsification: How does it relate to reproducibility? In J.-F. Morin, C. Olsson, & E. O. Atikcan (Eds.), *Key concepts in research methods*. Oxford, UK: Oxford University Press.

- Ebersole, C. R., Atherton, O. E., Belanger, A. L., Skulborstad, H. M., Allen, J. M., Banks, J. B., . . . Nosek, B. A. (2016). Many Labs 3: Evaluating participant pool quality across the academic semester via replication. *Journal of Experimental Social Psychology*, 67, 68–82. <http://dx.doi.org/10.1016/j.jesp.2015.10.012>
- Egger, M., Davey Smith, G., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal*, 315, 629–634. <http://dx.doi.org/10.1136/bmj.315.7109.629>
- Everett, J. A. C., & Earp, B. D. (2015). A tragedy of the (academic) commons: Interpreting the replication crisis in psychology as a social dilemma for early-career researchers. *Frontiers in Psychology*, 6, 1152. <http://dx.doi.org/10.3389/fpsyg.2015.01152>
- Everett, J. A. C., Pizarro, D. A., & Crockett, M. J. (2016). Inference of trustworthiness from intuitive moral judgments. *Journal of Experimental Psychology: General*, 145, 772–787. <http://dx.doi.org/10.1037/xge0000165>
- 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. <http://dx.doi.org/10.1037/0022-3514.69.6.1013>
- Festinger, L. (1957). *A theory of cognitive dissonance*. Stanford, CA: Stanford University Press.
- Festinger, L., & Carlsmith, J. M. (1959). Cognitive consequences of forced compliance. *The Journal of Abnormal and Social Psychology*, 58, 203–210. <http://dx.doi.org/10.1037/h0041593>
- Finkel, E. J., Eastwick, P. W., & Reis, H. T. (2015). Best research practices in psychology: Illustrating epistemological and pragmatic considerations with the case of relationship science. *Journal of Personality and Social Psychology*, 108, 275–297. <http://dx.doi.org/10.1037/pspi0000007>
- Finkel, E. J., Eastwick, P. W., & Reis, H. T. (2017). Replicability and other features of a high-quality science: Toward a balanced and empirical approach. *Journal of Personality and Social Psychology*, 113, 244–253. <http://dx.doi.org/10.1037/pspi0000075>
- Fisher, D. H. (1989). *Albion's seed: Four British folkways in America*. New York, NY: Oxford University Press.
- Fitz, N. S., Nadler, R., Manogaran, P., Chong, E. W. J., & Reiner, P. B. (2014). Public attitudes toward cognitive enhancement. *Neuroethics*, 7, 173–188. <http://dx.doi.org/10.1007/s12152-013-9190-z>
- Foot, P. (2001). *Natural goodness*. New York, NY: Oxford University Press.
- Forsell, E., Viganola, D., Pfeiffer, T., Almenberg, J., Wilson, B., Chen, Y., . . . Dreber, A. (2019). Predicting replication outcomes in the Many Labs 2 study. *Journal of Economic Psychology*, 75, 102117. <http://dx.doi.org/10.1016/j.joep.2018.10.009>
- Galinsky, A. D., Leonardelli, G. J., Okhuysen, G. A., & Mussweiler, T. (2005). Regulatory focus at the bargaining table: Promoting distributive and integrative success. *Personality and Social Psychology Bulletin*, 31, 1087–1098. <http://dx.doi.org/10.1177/0146167205276429>
- Galinsky, A. D., & Mussweiler, T. (2001). First offers as anchors: The role of perspective-taking and negotiator focus. *Journal of Personality and Social Psychology*, 81, 657–669. <http://dx.doi.org/10.1037/0022-3514.81.4.657>
- Galton, F. (1907). Vox populi. *Nature*, 75, 450–451. <http://dx.doi.org/10.1038/075450a0>
- 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. <http://dx.doi.org/10.1037/0033-2909.132.5.692>
- Gelman, A., & Stern, H. (2006). The difference between “significant” and “not significant” is not itself statistically significant. *The American Statistician*, 60, 328–331. <http://dx.doi.org/10.1198/000313006X152649>
- Giner-Sorolla, R. (2012). Science or art? How aesthetic standards grease the way through the publication bottleneck but undermine science. *Perspectives on Psychological Science*, 7, 562–571. <http://dx.doi.org/10.1177/1745691612457576>
- Greene, J. (2013). *Moral tribes: Emotion, reason, and the gap between us and them*. New York, NY: Penguin Press.
- Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., & Cohen, J. D. (2001). An fMRI Investigation of Emotional Engagement in Moral Judgment. *Science*, 293, 2105–2108.
- Greenland, S. (2017). Invited commentary: The need for cognitive science in methodology. *American Journal of Epidemiology*, 186, 639–645. <http://dx.doi.org/10.1093/aje/kwx259>
- Greenwald, A. G. (1975). Consequences of prejudice against the null hypothesis. *Psychological Bulletin*, 82, 1–20. <http://dx.doi.org/10.1037/h0076157>
- Greenwald, A. G., & Banaji, M. R. (1995). Implicit social cognition: Attitudes, self-esteem, and stereotypes. *Psychological Review*, 102, 4–27. <http://dx.doi.org/10.1037/0033-295X.102.1.4>
- Greenwald, A. G., McGhee, D. E., & Schwartz, J. L. K. (1998). Measuring individual differences in implicit cognition: The implicit association test. *Journal of Personality and Social Psychology*, 74, 1464–1480. <http://dx.doi.org/10.1037/0022-3514.74.6.1464>
- Groh, M., Krishnan, N., McKenzie, D., & Vishwanath, T. (2016). The impact of soft skill training on female youth employment: Evidence from a randomized experiment in Jordan. *IZA Journal of Labor and Development*, 5, 1–23. <http://dx.doi.org/10.1186/s40175-016-0055-9>
- Gronau, Q. F., Ly, A., & Wagenmakers, E.-J. (2017). *Informed Bayesian t-tests*. Manuscript submitted for publication. Retrieved from <https://arxiv.org/abs/1704.02479>
- Gronau, Q. F., Sarafoglou, A., Matzke, D., Ly, A., Boehm, U., Marsman, M., . . . Steingrover, H. (2017). A tutorial on bridge sampling. *Journal of Mathematical Psychology*, 81, 80–97. <http://dx.doi.org/10.1016/j.jmp.2017.09.005>
- Gronau, Q. F., Singmann, H., & Wagenmakers, E.-J. (2017). *bridgesampling: An R package for estimating normalizing constants*. Manuscript submitted for publication and uploaded to arXiv. Retrieved from <https://arxiv.org/abs/1710.08162>
- Gronau, Q. F., van Erp, S., Heck, D. W., Cesario, J., Jonas, K. J., & Wagenmakers, E. J. (2017). A Bayesian model-averaged meta-analysis of the power pose effect with informed and default priors: The case of felt power. *Comprehensive Results in Social Psychology*, 2, 123–138. <http://dx.doi.org/10.1080/23743603.2017.1326760>
- Gunia, B. C., Swaab, R. I., Sivanathan, N., & Galinsky, A. D. (2013). The remarkable robustness of the first-offer effect: Across culture, power, and issues. *Personality and Social Psychology Bulletin*, 39, 1547–1558. <http://dx.doi.org/10.1177/0146167213499236>
- Hahn, A., & Gawronski, B. (2019). Facing one's implicit biases: From awareness to acknowledgment. *Journal of Personality and Social Psychology*, 116, 769–794. <http://dx.doi.org/10.1037/pspi0000155>
- Hahn, A., Judd, C. M., Hirsh, H. K., & Blair, I. V. (2014). Awareness of implicit attitudes. *Journal of Experimental Psychology: General*, 143, 1369–1392. <http://dx.doi.org/10.1037/a0035028>
- Hendrick, C. (1990). Replications, strict replications, and conceptual replications: Are they important? *Journal of Social Behavior & Personality*, 5, 41–49.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, 33, 61–83. <http://dx.doi.org/10.1017/S0140525X0999152X>
- Higgins, J. P. T., Thompson, S. G., Deeks, J. J., & Altman, D. G. (2003). Measuring inconsistency in meta-analyses. *British Medical Journal*, 327, 557–560. <http://dx.doi.org/10.1136/bmj.327.7414.557>
- Huedo-Medina, T. B., Sánchez-Meca, J., Marín-Martínez, F., & Botella, J. (2006). Assessing heterogeneity in meta-analysis: Q statistic or I² index? *Psychological Methods*, 11, 193–206. <http://dx.doi.org/10.1037/1082-989X.11.2.193>
- Hüffmeier, J., Freund, P. A., Zerres, A., Backhaus, K., & Hertel, G. (2014). Being tough or being nice? A meta-analysis on the impact of hard-and

- softline strategies in distributive negotiations. *Journal of Management*, 40, 866–892.
- Hunter, J. E., & Schmidt, F. L. (2000). Fixed effects vs. random effects meta-analysis models: Implications for cumulative research knowledge. *International Journal of Selection and Assessment*, 8, 275–292. <http://dx.doi.org/10.1111/1468-2389.00156>
- Ibarra, H. (1995). Race, opportunity and diversity of social circles in managerial managers' networks. *Academy of Management Journal*, 38, 673–703.
- Ibarra, H. (1997). Paving an alternate route: Gender differences in network strategies for career development. *Social Psychology Quarterly*, 60, 91–102. <http://dx.doi.org/10.2307/2787014>
- Inglehart, R. (1997). *Modernization and postmodernization: Cultural, economic, and political change in 43 societies*. Princeton, NJ: Princeton University Press.
- Inglehart, R., & Welzel, C. (2005). *Modernization, cultural change, and democracy: The human development sequence*. Cambridge, MA: Cambridge University Press.
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS Medicine*, 2(8), e124. <http://dx.doi.org/10.1371/journal.pmed.0020124>
- Ioannidis, J. P. A., & Trikalinos, T. A. (2007). An exploratory test for an excess of significant findings. *Clinical Trials*, 4, 245–253. <http://dx.doi.org/10.1177/1740774507079441>
- JASP Team. (2018). JASP (Version 0.9.2.0) [Computer software]. Retrieved from <https://jasp-stats.org/>
- Jordan, J. J., Hoffman, M., Bloom, P., & Rand, D. G. (2016). Third-party punishment as a costly signal of trustworthiness. *Nature*, 530, 473–476. <http://dx.doi.org/10.1038/nature16981>
- Judd, C. M., Westfall, J., & Kenny, D. A. (2012). Treating stimuli as a random factor in social psychology: A new and comprehensive solution to a pervasive but largely ignored problem. *Journal of Personality and Social Psychology*, 103, 54–69. <http://dx.doi.org/10.1037/a0028347>
- Kahane, G. (2015). Sidetracked by trolleys: Why sacrificial moral dilemmas tell us little (or nothing) about utilitarian judgment. *Social Neuroscience*, 10, 551–560. <http://dx.doi.org/10.1080/17470919.2015.1023400>
- Kahneman, D., Diener, E., & Schwarz, N. (Eds.). (1999). *Well-being: The foundations of hedonic psychology*. New York, NY: Russell Sage Foundation.
- Kant, I. (1993). *Grounding of the metaphysics of morals* (3rd ed. Trans. J. W. Ellington). Indianapolis, IN: Hackett (Original work published 1775).
- Kirby, K. N., & Gerlanc, D. (2013). BootES: An R package for bootstrap confidence intervals on effect sizes. *Behavior Research Methods*, 45, 905–927. <http://dx.doi.org/10.3758/s13428-013-0330-5>
- Klein, R. A., Ratliff, K. A., Vianello, M., Adams, R. B., Jr., Bahník, Š., Bernstein, M. J., . . . Nosek, B. A. (2014). Investigating variation in replicability: A “many labs” replication project. *Social Psychology*, 45, 142–152. <http://dx.doi.org/10.1027/1864-9335/a000178>
- Kraut, R. (1979). Two conceptions of happiness. *The Philosophical Review*, 88, 167–197.
- Lakhani, K. R., Jeppesen, L. B., Lohse, P. A., & Panetta, J. A. (2007). *The value of openness in scientific problem solving*. Cambridge, MA: Division of Research, Harvard Business School.
- Lakhani, K. R., & Wolf, R. G. (2005). Why hackers do what they do: Understanding motivation and effort in free/open source software projects. In J. Feller, B. Fitzgerald, S. Hissam, & K. R. Lakhani (Eds.), *Perspectives on free and open source software* (pp. 3–21). Cambridge, MA: MIT Press.
- Landy, J. F., Walco, D. K., & Bartels, D. M. (2017). What's wrong with using steroids? Exploring whether and why people oppose the use of performance enhancing drugs. *Journal of Personality and Social Psychology*, 113, 377–392. <http://dx.doi.org/10.1037/pspa0000089>
- LeBel, E. P. (2015, October 13). *A list of successful and unsuccessful high-powered direct replications of social psychology findings*. Retrieved from <https://proveyourselfwrong.wordpress.com/2015/10/13/a-list-of-successful-and-unsuccessful-high-powered-direct-replications-of-social-psychology-findings/>
- LeBel, E. P., McCarthy, R. J., Earp, B. D., Elson, M., & Vanpaemel, W. (2018). A unified framework to quantify the credibility of scientific findings. *Advances in Methods and Practices in Psychological Science*, 1, 389–402. <http://dx.doi.org/10.1177/2515245918787489>
- Lieberman, M. D. (2013). *Social: Why our brains are wired to connect*. New York, NY: Crown Publishers.
- Lorenz, J., Rauhut, H., Schweitzer, F., & Helbing, D. (2011). How social influence can undermine the wisdom of crowd effect. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, 108, 9020–9025. <http://dx.doi.org/10.1073/pnas.1008636108>
- Loschelder, D. D., Swaab, R. I., Trötschel, R., & Galinsky, A. D. (2014). The first-mover disadvantage: The folly of revealing compatible preferences. *Psychological Science*, 25, 954–962. <http://dx.doi.org/10.1177/0956797613520168>
- Loschelder, D. D., Trötschel, R., Swaab, R. I., Friese, M., & Galinsky, A. D. (2016). The information-anchoring model of first offers: When moving first helps versus hurts negotiators. *Journal of Applied Psychology*, 101, 995–1012. <http://dx.doi.org/10.1037/apl0000096>
- Ma, D. S., Correll, J., & Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior Research Methods*, 47, 1122–1135. <http://dx.doi.org/10.3758/s13428-014-0532-5>
- Maaravi, Y., & Levy, A. (2017). When your anchor sinks your boat: Information asymmetry in distributive negotiations and the disadvantage of making the first offer. *Judgment and Decision Making*, 12, 420–429.
- Makel, M. C., Plucker, J. A., & Hegarty, B. (2012). Replications in psychology research how often do they really occur? *Perspectives on Psychological Science*, 7, 537–542. <http://dx.doi.org/10.1177/1745691612460688>
- McGuire, W. J. (1973). The yin and yang of progress in social psychology: Seven koan. *Journal of Personality and Social Psychology*, 26, 446–456. <http://dx.doi.org/10.1037/h0034345>
- McGuire, W. J. (1983). A contextualist theory of knowledge: Its implications for innovations and reform in psychological research. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 16, pp. 1–47). New York, NY: Academic Press. [http://dx.doi.org/10.1016/S0065-2601\(08\)60393-7](http://dx.doi.org/10.1016/S0065-2601(08)60393-7)
- McShane, B. B., & Gelman, A. (2017). Abandon statistical significance. *Nature*, 551, 558.
- McShane, B. B., Tackett, J. L., Böckenholt, U., & Gelman, A. (2019). Large scale replication projects in contemporary psychological research. *The American Statistician*, 73, 99–105. <http://dx.doi.org/10.1080/00031305.2018.1505655>
- Meng, X.-L., & Wong, W. H. (1996). Simulating ratios of normalizing constants via a simple identity: A theoretical exploration. *Statistica Sinica*, 6, 831–860.
- Mill, J. S. (2004). *Utilitarianism and other essays*. London, UK: Penguin Books. (Original work published 1861)
- Monin, B., & Oppenheimer, D. M. (2014). The limits of direct replications and the virtues of stimulus sampling [Commentary on Klein et al. (2014)]. *Social Psychology*, 45, 299–300. <http://dx.doi.org/10.1027/1864-9335/a000202>
- Monin, B., Pizarro, D., & Beer, J. (2007). Deciding vs. reacting: Conceptions of moral judgment and the reason-affect debate. *Review of General Psychology*, 11, 99–111. <http://dx.doi.org/10.1037/1089-2680.11.2.99>
- Morey, R. D., & Rouder, J. N. (2015). BayesFactor 0.9.111. Comprehensive R Archive Network. Retrieved from <http://cran.r-project.org/web/packages/BayesFactor/index.html>

- Morris, S. B., & DeShon, R. P. (2002). Combining effect size estimates in meta-analysis with repeated measures and independent-groups designs. *Psychological Methods*, 7, 105–125. <http://dx.doi.org/10.1037/1082-989X.7.1.105>
- Nickerson, R. S. (1998). Confirmation bias: A ubiquitous phenomenon in many guises. *Review of General Psychology*, 2, 175–220. <http://dx.doi.org/10.1037/1089-2680.2.2.175>
- Nosek, B. A., Ebersole, C. R., DeHaven, A. C., & Mellor, D. T. (2018). The preregistration revolution. *Proceedings of the National Academy of Sciences of the United States of America*, 115, 2600–2606. <http://dx.doi.org/10.1073/pnas.1708274114>
- Nosek, B. A., Spies, J. R., & Motyl, M. (2012). Scientific utopia: II. Restructuring incentives and practices to promote truth over publishability. *Perspectives on Psychological Science*, 7, 615–631. <http://dx.doi.org/10.1177/1745691612459058>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349, aac4716. <http://dx.doi.org/10.1126/science.aac4716>
- Pashler, H., & Harris, C. R. (2012). Is the replicability crisis overblown? Three arguments examined. *Perspectives on Psychological Science*, 7, 531–536. <http://dx.doi.org/10.1177/1745691612463401>
- Pfeiffer, T., Bertram, L., & Ioannidis, J. P. (2011). Quantifying selective reporting and the Proteus phenomenon for multiple datasets with similar bias. *PLoS ONE*, 6, e18362. <http://dx.doi.org/10.1371/journal.pone.0018362>
- Phillips, J., De Freitas, J., Mott, C., Gruber, J., & Knobe, J. (2017). True happiness: The role of morality in the folk concept of happiness. *Journal of Experimental Psychology: General*, 146, 165–181. <http://dx.doi.org/10.1037/xge0000252>
- Phillips, J., Nyholm, S., & Liao, S. (2014). The good in happiness. In J. Knobe, T. Lombrozo, & S. Nichols (Eds.), *Oxford studies in experimental philosophy* (Vol. 1, pp. 253–293). New York, NY: Oxford University Press.
- Poehlman, T. A. (2007). *Ideological inheritance: Implicit Puritanism in American moral cognition* (Doctoral dissertation). Yale University, New Haven, CT.
- R Core Team. (2018). R: A language and environment for statistical computing [Computer software manual]. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>
- Riley, R. D., Price, M. J., Jackson, D., Wardle, M., Gueyffier, F., Wang, J., . . . White, I. R. (2015). Multivariate meta-analysis using individual participant data. *Research Synthesis Methods*, 6, 157–174. <http://dx.doi.org/10.1002/jrsm.1129>
- Rosenthal, R. (1979). The “file drawer problem” and the tolerance for null results. *Psychological Bulletin*, 86, 638–641. <http://dx.doi.org/10.1037/0033-2909.86.3.638>
- Rosenthal, R., & DiMatteo, M. R. (2001). Meta-analysis: Recent developments in quantitative methods for literature reviews. *Annual Review of Psychology*, 52, 59–82. <http://dx.doi.org/10.1146/annurev.psych.52.1.59>
- Ryff, C. (1989). Happiness is everything, or is it? Explorations on the meaning of psychological well-being. *Journal of Personality and Social Psychology*, 57, 1069–1081. <http://dx.doi.org/10.1037/0022-3514.57.6.1069>
- Sanders, M., Mitchell, F., & Chonaire, A. N. (2015). *Just common sense? How well do experts and lay-people do at predicting the findings of behavioural science experiments*. Working paper.
- Sattler, S., Forlini, C., Racine, E., & Sauer, C. (2013). Impact of contextual factors and substance characteristics on perspectives toward cognitive enhancement. *PLoS ONE*, 8, e71452. <http://dx.doi.org/10.1371/journal.pone.0071452>
- Scheibehenne, B., Gronau, Q. F., Jamil, T., & Wagenmakers, E.-J. (2017). Fixed or random? A resolution through model-averaging. Reply to Carlsson, Schimmack, Williams, and Burkner (2017). *Psychological Science*, 28, 1698–1701. <http://dx.doi.org/10.1177/0956797617724426>
- Scheske, C., & Schnall, S. (2012). The ethics of “smart drugs”: Moral judgments about healthy people’s use of cognitive-enhancing drugs. *Basic and Applied Social Psychology*, 34, 508–515. <http://dx.doi.org/10.1080/01973533.2012.711692>
- Schimmack, U. (2012). The ironic effect of significant results on the credibility of multiple-study articles. *Psychological Methods*, 17, 551–566. <http://dx.doi.org/10.1037/a0029487>
- Schmidt, S. (2009). Shall we really do it again? The powerful concept of replication is neglected in the social sciences. *Review of General Psychology*, 13, 90–100. <http://dx.doi.org/10.1037/a0015108>
- Schweinsberg, M. (2013). *Starting high shrinks the pie*. Unpublished raw data, Department of Organizational Behavior, European School of Management and Technology Berlin, Berlin, Germany.
- Schweinsberg, M., Ku, G., Wang, C. S., & Pillutla, M. M. (2012). Starting high and ending with nothing: The role of anchors and power in negotiations. *Journal of Experimental Social Psychology*, 48, 226–231. <http://dx.doi.org/10.1016/j.jesp.2011.07.005>
- Schweinsberg, M., Madan, N., Vianello, M., Sommer, S. A., Jordan, J., Tierney, W., . . . Uhlmann, E. L. (2016). The pipeline project: Pre-publication independent replications of a single laboratory’s research pipeline. *Journal of Experimental Social Psychology*, 66, 55–67. <http://dx.doi.org/10.1016/j.jesp.2015.10.001>
- Schweinsberg, M., Viganola, D., Prasad, V., Dreber, A., Johannesson, M., Pfeiffer, T., . . . Uhlmann, E. L. (2019). *The pipeline project 2: Opening pre-publication independent replication to the world*. Manuscript submitted for publication.
- Silberzahn, R., & Uhlmann, E. L. (2015). Many hands make tight work: Crowdsourcing research can balance discussions, validate findings and better inform policy. *Nature*, 526, 189–191. <http://dx.doi.org/10.1038/526189a>
- Silberzahn, R., Uhlmann, E. L., Martin, D. P., Anselmi, P., Aust, F., Awtry, E., . . . Carlsson, R. (2018). Many analysts, one data set: Making transparent how variations in analytic choices affect results. *Advances in Methods and Practices in Psychological Science*, 1, 337–356. <http://dx.doi.org/10.1177/2515245917747646>
- Simons, D. J. (2014). The value of direct replication. *Perspectives on Psychological Science*, 9, 76–80. <http://dx.doi.org/10.1177/1745691613514755>
- Simonsohn, U., Nelson, L. D., & Simmons, J. P. (2014). P-curve: A key to the file-drawer. *Journal of Experimental Psychology: General*, 143, 534–547. <http://dx.doi.org/10.1037/a0033242>
- Simonsohn, U., Simmons, J., & Nelson, L. (2016). *Specification curve: Descriptive and inferential statistics for all plausible specifications*. Unpublished manuscript, Department of Behavioral Science, ESADE Business School, Barcelona, Spain.
- Singer, P. (2005). Ethics and intuitions. *The Journal of Ethics*, 9, 331–352. <http://dx.doi.org/10.1007/s10892-005-3508-y>
- Sowden, W., & Hall, M. (2015). *Exploring the relationship between morality and happiness*. Unpublished raw data, Department of Psychology, University of Michigan, Ann Arbor, MI.
- Stan Development Team. (2018). RStan: The R interface to Stan (R package version 2.17.3) [Computer software]. Retrieved from <http://mc-stan.org/>
- Steege, S., Tuerlinckx, F., Gelman, A., & Vanpaemel, W. (2016). Increasing transparency through a multiverse analysis. *Perspectives on Psychological Science*, 11, 702–712. <http://dx.doi.org/10.1177/1745691616658637>
- Sternberg, R. J. (2016). “Am I famous yet?” Judging scholarly merit in psychological science an introduction. *Perspectives on Psychological Science*, 11, 877–881. <http://dx.doi.org/10.1177/1745691616661777>
- Stewart, N., Ungemach, C., Harris, A. J. L., Bartels, D. M., Newell, B. R., Paolacci, G., & Chandler, J. (2015). The average laboratory samples a population of 7,300 Amazon Mechanical Turk Workers. *Judgment and Decision Making*, 10, 479–491.

- Stroebe, W., & Strack, F. (2014). The alleged crisis and the illusion of exact replication. *Perspectives on Psychological Science*, 9, 59–71. <http://dx.doi.org/10.1177/1745691613514450>
- Surowiecki, J. (2004). *The wisdom of crowds*. New York, NY: Random House.
- Tierney, W., Ebersole, C., Hardy, J., Chapman, H., Gantman, A., Vanaman, M., . . . Uhlmann, E. L. (2019). A creative destruction approach to replication. (Registered Report proposal under review)
- Tierney, W. T., Ebersole, C., Hardy, J. D., Gordon, M., Dreber, A., Johannesson, M., . . . Uhlmann, E. L. (2019). Replication ring for priming effects on judgments and behaviors. (Registered Report proposal in preparation)
- Tiokhin, L., Hackman, J., Munira, S., Jesmin, K., & Hruschka, D. (2019). Generalizability is not optional: Insights from a cross-cultural study of social discounting. *Royal Society Open Science*, 6, 181386. <http://dx.doi.org/10.1098/rsos.181386>
- Turiel, E. (1983). *The development of social knowledge: Morality and convention*. Cambridge, UK: Cambridge University Press.
- Turiel, E. (2002). *The culture of morality*. Cambridge, UK: Cambridge University Press.
- Tversky, A., & Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science*, 185, 1124–1131. <http://dx.doi.org/10.1126/science.185.4157.1124>
- Uhlmann, E. L., & Cunningham, W. A. (2000). *Awareness of automatic prejudice*. Unpublished raw data, Department of Organizational Behavior, INSEAD, Singapore.
- Uhlmann, E. L., Ebersole, C. R., Chartier, C. R., Errington, T. M., Kidwell, M. C., Lai, C. K., . . . Nosek, B. A. (2019). Scientific utopia III: Crowdsourcing science. *Perspectives on Psychological Science*, 14, 711–733. <http://dx.doi.org/10.1177/1745691619850561>
- Uhlmann, E. L., Pizarro, D. A., & Diermeier, D. (2015). A person-centered approach to moral judgment. *Perspectives on Psychological Science*, 10, 72–81. <http://dx.doi.org/10.1177/1745691614556679>
- Uhlmann, E. L., Poehlman, T. A., & Bargh, J. A. (2009). American moral exceptionalism. In J. T. Jost, A. C. Kay, & H. Thorisdottir (Eds.), *Social and psychological bases of ideology and system justification* (pp. 27–52). New York, NY: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780195320916.003.002>
- Uhlmann, E. L., Poehlman, T. A., & Nosek, B. A. (2012). Automatic associations: Personal attitudes or cultural knowledge? In J. Hanson (Ed.), *Ideology, psychology, and law* (pp. 228–260). Oxford, UK: Oxford University Press. <http://dx.doi.org/10.1093/acprof:oso/9780199737512.003.0009>
- Uhlmann, E. L., & Sanchez-Burks, J. (2014). The implicit legacy of American Protestantism. *Journal of Cross-Cultural Psychology*, 45, 992–1005. <http://dx.doi.org/10.1177/0022022114527344>
- Van Bavel, J. J., Mende-Siedlecki, P., Brady, W. J., & Reinero, D. A. (2016). Contextual sensitivity in scientific reproducibility. *PNAS Proceedings of the National Academy of Sciences of the United States of America*, 113, 6454–6459. <http://dx.doi.org/10.1073/pnas.1521897113>
- van den Bergh, D., van Doorn, J., Marsman, M., Draws, T., van Kesteren, E., Derks, K., . . . Wagenmakers, E.-J. (2019). *How to interpret the output of a Bayesian ANOVA in JASP*. Manuscript submitted for publication.
- van Erp, S., Verhagen, J., Grasman, R. P. P. P., & Wagenmakers, E.-J. (2017). Estimates of between-study heterogeneity for 705 meta-analyses reported in *Psychological Bulletin* from 1990–2013. *Journal of Open Psychology Data*, 5, 4. <http://dx.doi.org/10.5334/jopd.33>
- Vazire, S. (2017). Our obsession with eminence warps research. *Nature*, 547, 7. <http://dx.doi.org/10.1038/547007a>
- Verhagen, J., & Wagenmakers, E.-J. (2014). Bayesian tests to quantify the result of a replication attempt. *Journal of Experimental Psychology: General*, 143, 1457–1475. <http://dx.doi.org/10.1037/a0036731>
- Viechtbauer, W. (2010). Conducting meta-analyses in R with the metafor package. *Journal of Statistical Software*, 36, 1–48. <http://dx.doi.org/10.18637/jss.v036.i03>
- Wagenmakers, E.-J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., . . . Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin & Review*, 25, 58–76. <http://dx.doi.org/10.3758/s13423-017-1323-7>
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., . . . Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin & Review*, 25, 35–57. <http://dx.doi.org/10.3758/s13423-017-1343-3>
- Wagenmakers, E.-J., Wetzels, R., Borsboom, D., van der Maas, H. L. J., & Kievit, R. A. (2012). An agenda for purely confirmatory research. *Perspectives on Psychological Science*, 7, 632–638. <http://dx.doi.org/10.1177/1745691612463078>
- Waterman, A. S. (1993). Two conceptions of happiness: Contrasts of personal expressiveness (Eudaimonia) and hedonic enjoyment. *Journal of Personality and Social Psychology*, 64, 678–691. <http://dx.doi.org/10.1037/0022-3514.64.4.678>
- Waterman, A. S., Schwartz, S. J., Zamboanga, B. L., Ravert, R. D., Williams, M. K., Agocha, V. B., . . . Donnellan, B. (2010). The questionnaire for Eudaimonic well-being: Psychometric properties, demographic comparisons, and evidence of validity. *The Journal of Positive Psychology*, 5, 41–61. <http://dx.doi.org/10.1080/17439760903435208>
- Watson, D., Clark, L. A., & Tellegen, A. (1988). Development and validation of brief measures of positive and negative affect: The PANAS scales. *Journal of Personality and Social Psychology*, 54, 1063–1070. <http://dx.doi.org/10.1037/0022-3514.54.6.1063>
- Wells, G. L., & Windschitl, P. D. (1999). Stimulus sampling in social psychological experimentation. *Personality and Social Psychology Bulletin*, 25, 1115–1125. <http://dx.doi.org/10.1177/01461672992512005>
- Westfall, J., Judd, C. M., & Kenny, D. A. (2015). Replicating studies in which samples of participants respond to samples of stimuli. *Perspectives on Psychological Science*, 10, 390–399. <http://dx.doi.org/10.1177/1745691614564879>
- Wicherts, J. M., Veldkamp, C. L., Augusteijn, H. E., Bakker, M., van Aert, R. C., & van Assen, M. A. (2016). Degrees of freedom in planning, running, analyzing, and reporting psychological studies: A checklist to avoid p-hacking. *Frontiers in Psychology*, 7, 1832. <http://dx.doi.org/10.3389/fpsyg.2016.01832>
- Wilson, T. D., Lindsey, S., & Schooler, T. Y. (2000). A model of dual attitudes. *Psychological Review*, 107, 101–126. <http://dx.doi.org/10.1037/0033-295X.107.1.101>
- Wuchty, S., Jones, B. F., & Uzzi, B. (2007a). The increasing dominance of teams in production of knowledge. *Science*, 316, 1036–1039. <http://dx.doi.org/10.1126/science.1136099>
- Wuchty, S., Jones, B., & Uzzi, B. (2007b). Why do team authored papers get cited more? *Science*, 317, 1496–1498. <http://dx.doi.org/10.1126/science.317.5844.1496b>
- Zwaan, R. A., Etz, A., Lucas, R. E., & Donnellan, M. B. (2018). Improving social and behavioral science by making replication mainstream: A response to commentaries. *Behavioral and Brain Sciences*, 41, e157. <http://dx.doi.org/10.1017/S0140525X18000961>

(Appendix follows)

Appendix

The Crowdsourcing Hypothesis Tests Collaboration

Matúš Adamkovič, University of Prešov; Ravin Alaei, University of Toronto; Casper J. Albers, University of Groningen; Aurélien Allard, University of Paris VIII; Ian A. Anderson, INSEAD; Michael R. Andreychik, Fairfield University; Peter Babinčák, University of Prešov; Bradley J. Baker, University of Massachusetts; Gabriel Baník, University of Prešov; Ernest Baskin, Saint Joseph's University; Jozef Bavolar, Pavol Josef Šafárik University in Košice; Ruud M. W. J. Berkers, Max Planck Institute for Human Cognitive & Brain Sciences; Michał Bialek, Kozminski University; Joel Blanke, Stockholm School of Economics; Johannes Breuer, GESIS—Leibniz Institute for the Social Sciences; Ambra Brizi, Sapienza University of Rome; Stephanie E. V. Brown, Texas A&M University; Florian Brühlmann, University of Basel; Hendrik Bruns, University of Hamburg; Leigh Caldwell, Irrational Agency; Jean-François Campourcy, Université Clermont Auvergne; Eugene Y. Chan, Monash University; Yen-Ping Chang, Institute of Sociology, Academia Sinica, Taiwan; Benjamin Y. Cheung, University of British Columbia; Alycia Chin,* Public Company Accounting Oversight Board; Kit W. Cho, University of Houston—Downtown; Simon Columbus, Vrije Universiteit Amsterdam; Paul Conway, Florida State University; Conrad A. Corretti, The University of Texas at Dallas; Adam W. Craig, University of Kentucky; Paul G. Curran, Grand Valley State University; Alexander F. Danvers, University of Arizona; Ian G. J. Dawson, University of Southampton; Martin V. Day, Memorial University of Newfoundland; Erik Dietl, Loughborough University; Johannes T. Doerflinger, University of Konstanz; Alice Dominici, European University Institute; Vilnius Dranseika, Kaunas University of Technology and Vilnius University; Peter A. Edelsbrunner, ETH Zurich; John E. Edlund, Rochester Institute of Technology; Matthew Fisher, Southern Methodist University; Anna Fung, University of Washington; Oliver Genschow, University of Cologne; Timo Gnambs, Leibniz Institute for Educational Trajectories and Johannes Kepler University Linz; Matthew H. Goldberg, Yale University; Lorenz Graf-Vlachy, University of Passau; Andrew C. Hafenbrack, University of Washington; Sebastian Hafenbrädl, IESE Business School; Andree Hartanto, Singapore Management University; Patrick R. Heck, Geisinger Health System; Joseph P. Heffner, Brown University; Joseph Hilgard, Illinois State University; Felix Holzmeister, University of Innsbruck; Oleksandr V. Horchak, Instituto Universitário de Lisboa (ISCTE-IUL), CIS-IUL; Tina S.-T. Huang, University College London; Joachim Hüffmeier, TU Dortmund University; Sean Hughes and Ian Hussey, Ghent University; Roland Imhoff, Johannes Gutenberg University Mainz; Bastian Jaeger, Tilburg University; Konrad Jamro, University of Massachusetts Dartmouth; Samuel G. B. Johnson, University of Bath; Andrew Jones, University of Liverpool; Lucas Keller, University of Konstanz; Olga Kombeiz, Loughborough University; Lacy E. Krueger, Texas A&M University-Commerce; Anthony Lantian, Université Paris Nanterre; Jus-

tin P. Laplante, Clark University; Ljiljana B. Lazarevic, University of Belgrade; Jonathan Leclerc, John Molson School of Business, Concordia University; Nicole Legate, Illinois Institute of Technology; James M. Leonhardt, University of Nevada, Reno; Desmond W. Leung, Baruch College and The Graduate Center, City University of New York; Carmel A. Levitan, Occidental College; Hause Lin, University of Toronto; Qinglan Liu, Hubei University; Marco Tullio Liuzza, Magna Græcia University of Catanzaro; Kenneth D. Locke, University of Idaho; Albert L. Ly, Loma Linda University; Melanie MacEacheron, University of Western Ontario; Christopher R. Madan, University of Nottingham; Harry Manley, Chulalongkorn University; Silvia Mari, University of Milano - Bicocca; Marcel Martončík, University of Prešov; Scott L. McLean, Walden University; Jonathon McPhetres, Massachusetts Institute of Technology and University of Regina; Brett G. Mercier, University of California, Irvine; Corinna Michels, University of Cologne; Michael C. Mullarkey, University of Texas-Austin; Erica D. Musser, Florida International University; Ladislav Nalborczyk, Univ. Grenoble Alpes and Ghent University; Gustav Nilsson, Karolinska Institutet and Stockholm University; Nicholas G. Otis, University of California, Berkeley; Sarah M. G. Otner, Imperial College Business School; Philipp E. Otto, European University Viadrina; Oscar Oviedo-Trespalacios, Queensland University of Technology and Universidad del Norte; Mariola Paruzel-Czachura, University of Silesia; Francesco Pellegrini, Università degli Studi di Padova; Vitor M. D. Pereira, Universidade de Lisboa, Alameda da Universidade; Hannah Perfecto, Washington University in St. Louis; Gerit Pfuhl, UiT The Arctic University of Norway; Mark H. Phillips, Abilene Christian University; Ori Plonsky, Duke University; Maura Pozzi, Università Cattolica del Sacro Cuore; Danka B. Purić, University of Belgrade; Brett Raymond-Barker, University of Roehampton; David E. Redman, Pacific Lutheran University; Caleb J. Reynolds, Florida State University; Ivan Ropovik, University of Prešov; Lukas Röseler, Harz University of Applied Sciences and University of Bamberg; Janna K. Ruessmann, University of Cologne; William H. Ryan, University of California, Berkeley; Nika Sablaturova, Masaryk University; Kurt J. Schuepfer, Miami University; Astrid Schütz, University of Bamberg; Miroslav Sirota, University of Essex; Matthias Stefan, University of Innsbruck; Eric L. Stocks, University of Texas at Tyler; Garrett L. Strosser, Southern Utah University; Jordan W. Suchow, Stevens Institute of Technology; Anna Szabelska, Queen's University Belfast; Kian Siong Tey, INSEAD; Leonid Tiokhin, Eindhoven University of Technology; Jais Troian, Istanbul Bilgi University; Till Utesch, University of Münster; Alejandro Vásquez-Echeverría, University of the Republic; Leigh Ann Vaughn, Ithaca College; Mark Verschoor, University of Groningen; Bettina von Helversen, University of Bremen; Pascal Wallisch, New York University; Sophia C. Weissgerber, University of Kassel; Aaron L. Wichman, Western Kentucky University;

(Appendix continues)

Jan K. Woike, Max Planck Institute for Human Development and DIW, Berlin, Germany; Iris Žeželj, University of Belgrade; Janis H. Zickfeld, University of Oslo and University of Mannheim; Yeonsin Ahn, Philippe F. Blaettchen, Xi Kang, Yoo Jin Lee, Philip M. Parker, Paul A. Parker, Jamie S. Song, May-Anne Very, and Lynn Wong, INSEAD.

The first through 135th members of the Crowdsourcing Hypothesis Tests Collaboration lent their expertise as evaluators of study quality and forecasters. The 136th through 144th members of the Crowdsourcing Hypothesis Tests Collaboration lent their expertise as evaluators of study quality and forecasters in a pilot version of the Forecasting Study. In addition, all members of the Crowd-

sourcing Hypothesis Tests Collaboration, along with the authors whose full names are listed on the first page, revised and approved the final version of this article.

* The Public Company Accounting Oversight Board, as a matter of policy, disclaims responsibility for any private publication or statement by any of its economic research fellows, consultants, or employees.

Received June 29, 2018

Revision received October 24, 2019

Accepted October 29, 2019 ■