

# Moderation of Classic Social Psychological Effects by Demographics in the U.S. Adult Population: New Opportunities for Theoretical Advancement

David S. Yeager  
University of Texas at Austin

Jon A. Krosnick  
Stanford University

Penny S. Visser  
University of Chicago

Allyson L. Holbrook  
University of Illinois at Chicago

Alex M. Tahk  
University of Wisconsin-Madison

For decades, social psychologists have collected data primarily from college undergraduates and, recently, from haphazard samples of adults. Yet researchers have routinely presumed that thus observed treatment effects characterize “people” in general. Tests of seven highly cited social psychological phenomena (two involving opinion change resulting from social influence and five involving the use of heuristics in social judgments) using data collected from randomly sampled, representative groups of American adults documented generalizability of the six phenomena that have been replicated previously with undergraduate samples. The 1 phenomenon (a cross-over interaction revealing an ease of retrieval effect) that has not been replicated successfully previously in undergraduate samples was also not observed here. However, the observed effect sizes for the replicated phenomena were notably smaller on average than the meta-analytic effect sizes documented by past studies of college students. Furthermore, the phenomena were strongest among participants with the demographic characteristics of the college students who typically provided data for past published studies, even after correcting for publication bias in past studies using a new method, called the *behaviorally-informed file-drawer adjustment*. The six successful replications suggest that phenomena identified in traditional laboratory research also appear as expected in representative samples but more weakly, so observed effect sizes should be generalized with caution. The evidence of demographic moderators suggests interesting opportunities for future research to better understand the mechanisms of the effects and their limiting conditions.

**Keywords:** replication, representative sampling, heuristics and biases, persuasion, conformity

**Supplemental materials:** <http://dx.doi.org/10.1037/pspa0000171.supp>

Although social psychology “generally seeks principles to describe social behavior that hold across persons” (Reis & Gosling, 2010, p. 85; Cook & Groom, 2004), the field’s body of evidence collected during the last 70 years is predominantly from studies of

a select subgroup of the population: college students enrolled in psychology courses who completed experiments to fulfill course requirements. Today, the vast majority of evidence documenting what might be called “classic” findings that were discovered

This article was published Online First August 29, 2019.

David S. Yeager, Department of Psychology and Population Research Center, University of Texas at Austin; Jon A. Krosnick, Department of Communications, Political Science, and Psychology, Stanford University; Penny S. Visser, Department of Psychology, University of Chicago; Allyson L. Holbrook, Department of Public Administration and Psychology, University of Illinois at Chicago; Alex M. Tahk, Department of Political Science, University of Wisconsin-Madison.

Penny S. Visser is now an Independent scholar in Chicago, Illinois.

Jon A. Krosnick is University Fellow at Resources for the Future. We thank all the people who provided feedback on previous drafts of the manuscript, too numerous to name. Writing of the manuscript was sup-

ported in part by the William T. Grant Foundation (PI: David S. Yeager), and the National Institute of Child Health and Human Development (Grant 10.13039/100000071 R01HD084772-01, PI: David S. Yeager; and Grant P2C-HD042849, to the Population Research Center at The University of Texas at Austin). The data for this article are available at [osf.io/s9zax](https://osf.io/s9zax)

Correspondence concerning this article should be addressed to David S. Yeager, Department of Psychology and Population Research Center, University of Texas at Austin, 1 University Station #A8000, Austin, TX 78712 or to Jon A. Krosnick, Department of Communications, Political Science, and Psychology, Stanford University, 432 McClatchy Hall, Stanford, CA 94305. E-mail: [dyeager@utexas.edu](mailto:dyeager@utexas.edu) or [krosnick@stanford.edu](mailto:krosnick@stanford.edu)

decades ago, serve as foundational pillars for the field, are widely presumed to be true today, and are routinely discussed in textbooks and classrooms is from these studies of so-called “haphazard samples” (Visser, Krosnick, & Lavrakas, 2000) of college students. This continues to be true even though studies documenting new phenomena are increasingly based on haphazard samples of adults who were not scientifically selected from a defined population and instead volunteered to participate in online research (e.g., Amazon.com’s Mechanical Turk workers; Buhrmester, Kwang, & Gosling, 2011).

According to various scholars, many effects that have been of interest to social psychologists may occur more powerfully or perhaps even exclusively among people with the characteristics of college students, and may appear more weakly or not at all among others (Henrich, Heine, & Norenzayan, 2010; Sears, 1986; Wells, 1993). For example, Sears (1986) proposed that the narrow age range, high educational levels, and other demographic characteristics typical of college students make them different from other adults in ways that may limit the generalizability of findings (Henry, 2008; Van Lange, Schippers, & Balliet, 2011). Likewise, Wells (1993) said that “students are not typical” because of their restricted age range and educational levels and that ignoring these uniquenesses “place[s] student-based conclusions at substantial risk” (pp. 491–492). Petty and Cacioppo (1996) responded to these concerns by noting that once a phenomenon had been demonstrated in haphazard samples, its generalizability to representative samples of adults can be ascertained. But such assessment has rarely if ever been done regarding most classic findings.

This article takes on Petty and Cacioppo’s challenge, gauging whether a set of classic social psychological phenomena previously documented in many studies of college students also appear with comparable strength in representative samples of American adults who were truly randomly selected from the population. Random sampling is implemented by assuring that every member of the population has a known, nonzero probability of being selected, and extensive efforts must be made to elicit participation by all of those sampled individuals. This methodology is the bedrock of survey research that has yielded many of the most important findings in sociology, political science, economics, and other social science disciplines. Even today, despite response rates dropping progressively in recent decades (Keeter, Kennedy, Dimock, Best, & Craighill, 2006), true random sampling continues to yield strikingly accurate measurements of populations in surveys (MacInnis, Krosnick, Ho, & Cho, 2018; Yeager et al., 2011).<sup>1</sup>

A few studies have explored the comparability of findings obtained from nonprobability haphazard samples of adults with the findings obtained from college student participants (e.g., Peterson, 2001). For example, an umbrella review of 30 meta-analyses of psychological effects, many of which were classics in the social psychology literature, found that effect sizes generated with college students often differed in terms of magnitude and even direction from those found with haphazard samples of American adults (Peterson, 2001). Many relations appeared in one type of sample and not in the other. Likewise, behavioral economics and cognitive psychology experiments have sometimes produced stronger support for researchers’ hypotheses among student participants than among haphazard samples of nonstudent adults living outside the United States (see Henrich et al., 2010, for a review). However, some studies of framing, attention, perception, and decision-

making have found similar experimental effects with haphazard groups of adult participants and with college students (Berinsky, Huber, & Lenz, 2012; Crump, McDonnell, & Gureckis, 2013; Goodman, Cryder, & Cheema, 2013). This mixed evidence raises questions about whether canonical studies in social psychology would have yielded the same effects and effect sizes if conducted with representative, general public samples.

## The Present Research

The focus of this article is on prominent social psychological phenomena. In addition to gauging their generalizability to representative samples, this investigation explored moderation of effect sizes by the sorts of demographics that define typical college students. Such evidence of demographic moderation is helpful not only for understanding generalization (or lack thereof) but also, as we will illustrate, has useful implications for basic theory development. Understanding among whom an effect is strongest can affirm or challenge presumed psychological mechanisms of effects. The results of the present investigation were often surprising in this regard and justify rethinking widespread assumptions about the processes responsible for some phenomena.

Specifically, we explored seven social psychological phenomena in two arenas: (a) opinion change resulting from conformity or persuasion (two effects), and (b) heuristics and biases in social judgment (five effects). The two phenomena involving opinion change resulting from social influence were:

1. Conformity to a simply presented descriptive norm (Asch, 1952; Cialdini, 2003; Sherif, 1936). In many past studies, participants learned about the proportion of other people who held a particular opinion or who performed a particular behavior. In some studies, other people were portrayed as unanimous, and in other, more realistic studies, other people were portrayed as manifesting a nonunanimous, majority opinion. Our study examined the impact of the latter.
2. The effect of a content-laden persuasive message on attitudes as moderated by argument quality and need for cognition (e.g., Cacioppo, Petty, & Morris, 1983). In hundreds of past studies, participants were randomly assigned to read or hear a persuasive message containing either strong or weak arguments. The impact of this argument quality manipulation on amount of attitude change in the direction of the message was larger among people higher in need for cognition, because these people thought more carefully and more effectively recognized the quality of the arguments. We examined this two-way interaction.

Five additional phenomena we examined involved the use of heuristics in social judgment:

<sup>1</sup> The appearance of inaccuracy of predictions of election outcomes based on preelection polls in the United States (e.g., Edwards-Levy, 2017; Newkirk, 2016) and other countries (e.g., Israel, the United Kingdom; Hanretty, 2016; King & Sobelman, 2015) is due to the use of haphazard samples rather than random samples in the vast majority of such polls (McLean, Krosnick, & Tahnk, 2018).

3. Base-rate underutilization (using the “lawyer/engineer” problem; Kahneman & Tversky, 1973). Participants were asked to make a judgment about a person after being given base-rate information and individuating information about the person. In past studies, participants mostly ignored the base rate when making the judgment about the person and were influenced by individuating information.
4. The conjunction fallacy (using the “Linda” problem; Tversky & Kahneman, 1983). In past studies, when judging whether a person is more likely to belong to a category defined by one characteristic (e.g., a bank teller) or defined by the conjunction of two characteristics (e.g., a feminist and a bank teller), many participants mistakenly chose the second instead of the first, because of the resemblance of the added attribute (feminist) to other characteristics of the person.
5. Underappreciation of the law of large numbers (using the “hospital” problem; Tversky & Kahneman, 1974). When asked which of two samples of people (one large, the other small) was more likely to accurately reflect the characteristics of a population to which they belong, many participants failed to recognize that the larger sample was more likely to yield accurate results.
6. The false consensus effect (e.g., Ross, Greene, & House, 1977). This has routinely been demonstrated by a positive correlation between participants’ reports of their own opinion and their estimates of the prevalence of that opinion among other people.
7. The effect of “ease of retrieval” on self-perceptions (e.g., Schwarz et al., 1991). This has been demonstrated by asking participants to retrieve either a few instances in which they performed a behavior or many such instances, and then to rate the extent to which a relevant trait describes them. In past studies, more experimentally induced retrieval difficulty was associated with lower ratings of the degree to which the relevant trait described the participant.

These seven phenomena do not constitute a random sample of all social psychological phenomena that we could have investigated. In fact, drawing such a sample seems daunting, if not impossible, because doing so would require first defining a population of such phenomena, yet there is no obviously optimal way to do so. Therefore, it is best to think of this investigation as what might be called “fixed effects” of phenomena, rather than “random effects.” That is, the demonstrations reported below should not be presumed to generalize to any larger population of phenomena. Rather, these demonstrations are just that: demonstrations that should not yet be considered generalizable.

The phenomena we investigated seem well-suited to such an agenda, because they have been the focus of extensive research in the past and have been tremendously visible and impactful. For example, conformity, need for cognition moderating the effect of argument quality in persuasion, the false consensus effect, and the

ease of retrieval effect have been the subjects of meta-analyses (Bond & Smith, 1996; Cacioppo, Petty, Feinstein, & Jarvis, 1996; Mullen et al., 1985; Weingarten & Hutchinson, 2018), and collectively these have included the results of more than 800 effect sizes. In addition, large numbers of publications have cited each of these phenomena (see the [online supplemental material](#) for the counts).

Furthermore, some of these phenomena have been of interest to previous investigators who wished to explore moderation by demographics. For example, research on heuristics and biases has relied on “dual systems” theories to explain why individuals make reasoning errors (Kahneman, 2003; Stanovich & West, 1998; also called “thinking fast, thinking slow” in the popular press; Kahneman, 2011). Cognitive skills equip people to avoid errors when making judgments, so less skilled people, such as individuals with lower educational attainment (Brinch & Galloway, 2012; Ceci, 1991; Rietveld et al., 2014), may be more likely to use heuristics and manifest biases when making social judgments. Furthermore, social power may cause individuals to rely more on their “gut feelings” and respond more heuristically (Weick & Guinote, 2008), and so characteristics that are known to be associated with social power in the United States—such as being in the racial majority group, being male, and being wealthy (Fiske, 2010)—might predict greater heuristic responding.

Likewise, individuals’ tendencies to change their attitudes might depend on the strength of their self-concepts (which may be related to their age; see Erikson, 1968; Sears, 1986; Visser & Krosnick, 1998), individuals’ abilities to detect and respond to subtleties in their circumstances (which may be reflected in their education levels; Cacioppo, Petty, Kao, & Rodriguez, 1986; Eagly & Warren, 1976), by feelings of social power that make them resistant to attitude-change attempts (which may be proxied by age, income, or race/ethnicity, Eaton, Visser, Krosnick, & Anand, 2009; see Fiske, 2010), or by people’s cultural interdependence/collectivism (which may differ between the Midwest and other regions in the United States; Plaut, Markus, & Lachman, 2002; also see Bond & Smith, 1996; Henrich et al., 2010). Seeking to replicate the seven phenomena in representative samples and testing for moderation by demographic characteristics is therefore informative for theories of these effects’ mechanisms.

## Method

### Data

Seven studies were conducted via the Internet by the firm then called *Knowledge Networks* (now called *GfK Custom Research*) with two national random samples of English-speaking American adults ( $n = 1,338$  and  $2,132$ ). The persuasion and conformity studies were conducted with Sample 1, and the five social reasoning studies were conducted with Sample 2. These were the only classic phenomena we attempted to replicate in these studies.

Probability sampling to build the samples of participants for these studies began with Knowledge Networks recruiting people to join a “panel” of individuals (called the *KnowledgePanel*) who consented to complete online survey questionnaires regularly (about twice per week).<sup>2</sup> Knowledge Networks did so using list-assisted random digit dialing (RDD) for telephone interviewing in three stages, a

<sup>2</sup> An extended version of this description of the sampling methods appears in the [online supplemental material](#).

procedure that saves money and increases efficiency by avoiding interviewers calling telephone numbers that were not associated with working residential phones.

Before calling began, Knowledge Networks mailed “advance letters” to as many selected households as possible. Doing so required obtaining mailing addresses associated with the sampled telephone numbers. The set of generated telephone numbers was submitted to a commercial company that used published telephone directories (and possibly other information sources) to obtain a mailing address for as many phone numbers as possible (which was usually about 50%). Then advance letters were mailed.

All telephone numbers were then called by recruiters up to 15 times (if no one ever answered the phone) or up to 25 times (if someone ever answered the phone) to talk with an adult and invite the members of the household to join the panel. Any households that initially declined to join the panel were recontacted and encouraged to reconsider their decisions.

In 2002, when our surveys were conducted, many households lacked Internet access, and excluding them from the panel would have introduced bias. So Knowledge Networks provided free Internet access and an Internet connection device (Web TV; [Wikipedia Contributors, 2018](#)) to households that needed such access in order to complete online questionnaires. Staff members were available to provide assistance to panel members who were unfamiliar with the Internet or the equipment, to assist in setting up the equipment and to assist during the process of questionnaire completion.

For each of our surveys, a separate random sample of panelists was drawn, with unequal probabilities of selection within strata defined by age, gender, race, ethnicity, and region of residence, so that the participating sample matched the demographic characteristics of the nation. Invitations to complete our questionnaires were sent to selected panelists via mailed paper letters and via emails. Furthermore, e-mail reminders were sent to unresponsive panelists, and telephone calls were made to them if they remained unresponsive.

The Knowledge Networks method has yielded highly accurate measurements of the U.S. adult population. One study ([Yeager et al., 2011](#)) examined surveys conducted by Knowledge Networks during 2004 and found that the characteristics of the samples were extremely similar to those of the population, with an average absolute error of 3.4 percentage points, when estimating the proportion of the population with a particular characteristic, relative to benchmarks obtained from official government data sources (e.g., marital status, household income, number of adults living in the home, number of bedrooms in the home, employment status, smoking, drinking alcohol, possession of a passport, possession of a driver’s license). Similar accuracy was observed recently by [MacInnis et al. \(2018\)](#).

## Procedures

Immediately after all survey participants were recruited into the KnowledgePanel, they reported their sex, age, race, household income, education, and U.S. Census region of residence. During later waves of data collection, participants read information and answered questions in ways that matched or approximated those used in the canonical studies documenting each phenomenon that we examined (for descriptions of the procedures used, see the [online supplemental material](#)).<sup>3</sup>

**Conformity.** Participants read about the results of a national survey measuring American public opinion on an issue of government

policy. Half of the participants (selected randomly) read that most people favored the policy, and the other half of the participants read that most people opposed the policy. All participants then reported their own attitudes toward the policy. Conformity was gauged by the impact of the descriptive norm on people’s own opinions.

**Persuasion.** Participants reported their opinions toward capital punishment and answered questions measuring need for cognition. One week later, people who were initially neutral or positive toward capital punishment read either strong or weak arguments against it. (A separate nationally representative sample of American adults rated the strong arguments as significantly stronger than the weak arguments, as expected, as described in the [online supplemental material](#).) Then the participants reported their own opinions again. This design permitted assessing the impact of argument quality on persuasion and moderation of that effect by need for cognition.

**Base rate underutilization.** Participants were randomly assigned to be told that either 30% or 70% (randomly assigned) of a group of men were engineers and the rest were lawyers. Participants then read a description of one of the men who had been randomly selected from the set. The description sounded lawyer-like or engineer-like (randomly assigned). Participants then estimated the probability that he was an engineer. The effects of the base rate and of the individuating information on probability estimates were calculated.<sup>4</sup>

**Law of large numbers.** Participants were asked to judge which of two hospitals—one that birthed many babies each day and the other that birthed fewer babies daily—had more days on which more than 60% of the babies born were boys. Failure to select the hospital that birthed fewer babies was treated as a reasoning error, because smaller samples deviate more from population parameters.

**Conjunction fallacy.** Participants were asked to read a description of a woman who matched the stereotype of a feminist and ranked a series of statements about her in terms of their likelihood of being true, including one stating that she was a feminist, one stating that she was a bank teller, and one stating that she was both a feminist and a bank teller. Ranking the latter statement as more likely than her being a bank teller was treated as a reasoning error, because the conjunction of any two events should be less likely than either of the two events alone.<sup>5</sup>

<sup>3</sup> For most of the studies we conducted, we used experimental materials that were identical or nearly identical to those used in past studies (lawyer/engineer, conjunction fallacy, law of large numbers, ease of retrieval). For other studies, procedures were the same, but the attitude object examined was adapted to be suitable for the American public (conformity to a descriptive norm, persuasion, and the false consensus effect). Supplemental experiments conducted with additional nationally representative samples documented that the manipulations evoked the intended psychological processes as intended (see the [online supplemental material](#)).

<sup>4</sup> Participants were randomly assigned to see either the base rate first or the individuating information, but this had no effect on our conclusions (see the [online supplemental material](#)).

<sup>5</sup> A randomly selected half of the participants rated the likelihood that each statement about the woman was true, rather than ranking the statements. This variation in measurement approach did not alter any conclusions about the relations of demographics with the propensity to commit the conjunction fallacy error (see the [online supplemental material](#)). To produce effect size estimates that were directly comparable to those in the canonical study, which asked participants to rank the statements, we only used data from the participants who ranked the statements.

**False consensus effect.** Participants were asked whether they favored or opposed a government policy and reported the percent of American adults whom they thought favored the policy. The false consensus effect was gauged by assessing the association between participants' own attitudes and their perceptions of others' attitudes.

**Ease of retrieval.** Participants reported either six or 12 instances (randomly assigned) when they behaved assertively or unassertively (randomly assigned) in the past. Then, participants rated how assertive or unassertive (randomly assigned) they were. The ease of retrieval effect was gauged by testing whether ratings of assertiveness were higher after retrieving six instances of being assertive than after retrieving 12 instances of being assertive and higher after retrieving 12 instances of being unassertive than after retrieving six instances of being unassertive. That is, we tested for the canonical crossover interaction of the valence of instances retrieved from memory and the number of retrieved instances.

## Analytic Method

**Main effects and demographic moderation.** To gauge effect sizes in the full sample and to test for demographic moderation of those effect sizes, we used semiparametric generalized additive models (GAMs) with natural cubic splines (see, e.g., Andersen, 2009; Keele, 2008). Rather than imposing particular functional forms on the operation of continuous variables (e.g., linear or quadratic), GAMs use flexible regression curves (i.e., splines) across values of continuous predictors (Keele, 2008) to discover the best-fitting functional form but include a penalty to avoid overfitting (for an extensive discussion, see the [online supplemental material](#)). We chose GAMs to minimize arbitrary decision-making about the likely functional form of a continuous variable's impact and reduce the possibility that one kind of "researcher degree of freedom" could cause irreplicable results (Simmons, Nelson, & Simonsohn, 2011; also see Feller & Holmes, 2009).

Multicollinearity among the demographic moderators could cause problems in the estimation of the interaction terms' standard errors. However, the demographic moderators were not strongly correlated with one another—the strongest correlation was between education and income,  $r = .37$ , and the rest were far smaller. Therefore, multicollinearity is not likely to have caused imprecision in estimates of the interaction terms.

**Effect sizes from past studies.** Effect sizes from the present study were compared to effect sizes from the canonical study that most prominently showed the effect (usually the first study published), and from meta-analyses of subsequent studies of college students. If an effect had been examined in a published meta-analysis or systematic review, we used the effect size estimates from that analysis (Bond & Smith, 1996; Cacioppo et al., 1996; Hertwig & Chase, 1998; Mullen et al., 1985). When no such papers had been published, we conducted our own meta-analysis (see the [online supplemental material](#)). For the ease of retrieval effect, we used a published meta-analysis (Weingarten & Hutchinson, 2018) to identify studies that examined the crossover interaction, which was the primary effect in the original paper, and meta-analyzed those effects.

For conformity to a descriptive norm, the canonical study results reported here are from the three nonunanimous majority studies reported by Asch (1952). We did so because our replication involved information about a descriptive norm that was nonunanimous. Bond and Smith's (1996) meta-analysis of subsequent conformity studies

only examined studies that used a unanimous majority and therefore presumably yielded a stronger effect size than would be observed with nonunanimous majorities. Therefore, we adjusted Bond and Smith's (1996) meta-analytic average effect size for unanimous majority studies ( $d = .92$ ) by multiplying it by the ratio of the effect sizes from Asch's (1952) nonunanimous majority studies to the effect sizes from Asch's (1952) unanimous majority studies (which was .27).

**Simulated effect for a hypothetical group of participants with the characteristics of "college students."** We sought to gauge what the effect sizes would have been if the national survey participants had the demographic profile of college student participants in past social psychology studies.<sup>6</sup> We did so by following statistical recommendations for estimating conditional average treatment effects (Feller & Holmes, 2009) using the parameters of GAMs generated using the national survey data.

The demographics of college student participants were gauged using data from two sources. The first is the raw dataset created by Gosling, Vazire, Srivastava, and John (2004) documenting the ages, genders, races, and regions of the participants in studies described in *JSPS* in 2002. Using those data, we calculated the distributions of those demographics for participants.

To document the distributions of educational attainment and total family income that college student study participants would eventually earn in the year 2002 (the year of collection of the data reported here), we relied on data collected from the so-called "1979 cohort" of the National Longitudinal Survey of Youth, a large representative national sample of Americans who were ages 14 to 22 in 1979 (Bureau of Labor Statistics, U.S. Department of Labor, 2014). We chose this cohort because most of the studies we conducted were originally conducted in the 1970s and early 1980s. Among National Longitudinal Survey of Youth participants who attended a four-year college in the late 1970s and early 1980s, we computed the distributions of their educational attainment and total family income in 2002. Data from these two sources were then used to simulate a new dataset of individual observations (one per hypothetical study participant) in which the distributions of demographics matched the estimated distributions of the characteristics of typical college student participants. These are also called "synthetic observations" in related nonparametric approaches to heterogeneous treatment effects (Green & Kern, 2012).

This simulated dataset was then fed into the GAMs to predict  $\hat{y}_i$  for each hypothetical individual, separately for each of the effects that we studied. This method is necessary because GAMs involve nonparametric estimation, so generating predicted values ( $\hat{y}$ ) for a given value of a moderator cannot be accomplished via the sort of simple arithmetic that can be done with the results of ordinary least squares regressions. Using the simulated dataset described above, it was possible to aggregate the predicted values for all hypothetical individuals in the full sample, or separately for individuals in different experimental conditions, to yield estimates of effect sizes for each effect.

The equation used is

$$\widehat{y}_{i,a} = f_{1,a}(Age_i) + f_{2,a}(Education_i) + f_{3,a}(Income_i) + b_{1,a}(White_i) + b_{2,a}(Female_i) + b_{3,a}(Midwest_i) \quad (1)$$

<sup>6</sup> We assume that the demographic characteristics of study participants were similar across recent decades, but we know of no available data with which to test this assumption.

where, for participant  $i$ ,  $a$  is the value of the condition variable (for reasoning error studies,  $a = \{1\}$ ; in a two-cell design,  $a = \{0, 1\}$ ; in a four-cell,  $a = \{1, 2, 3, 4\}$ ),  $f_{1,a}$ ,  $f_{2,a}$ , and  $f_{3,a}$  are nonparametric functions estimated for each condition by the GAM (via thin plate regression splines),  $b_{1,a}$ ,  $b_{2,a}$ , and  $b_{3,a}$  are standard parametric regression coefficients estimated separately for each condition. We used the parameters from Equation (1) to calculate  $\hat{y}_i$  for a dataset of participants with length  $n$  (the number of observations for a given value of  $a$ ) and then computed the mean  $\hat{y}_{i,a}$  for each value of  $a$ . The simulated unstandardized effect sizes (ES) are:

- For probabilistic error studies:  $\widehat{ES} = \frac{\sum_{i=1}^n \hat{y}_i}{n}$ .
- For two condition studies:  $\widehat{ES} = \frac{\sum_{i=1}^{n_{a=0}} \hat{y}_{i,a=0}}{n_{a=0}} - \frac{\sum_{i=1}^{n_{a=1}} \hat{y}_{i,a=1}}{n_{a=1}}$
- For four condition studies:  $\widehat{ES} = \left( \frac{\sum_{i=1}^{n_{a=0}} \hat{y}_{i,a=0}}{n_{a=0}} - \frac{\sum_{i=1}^{n_{a=1}} \hat{y}_{i,a=1}}{n_{a=1}} \right) - \left( \frac{\sum_{i=1}^{n_{a=3}} \hat{y}_{i,a=3}}{n_{a=3}} - \frac{\sum_{i=1}^{n_{a=4}} \hat{y}_{i,a=4}}{n_{a=4}} \right)$

Then, we estimated standardized effect sizes (Cohen’s  $d$ ) by dividing the unstandardized estimates by the pooled standard deviations from the data.

The generalizability of the simulation of effect sizes is based upon two fundamental assumptions, as described by Allcott (2015). The first is external unconfoundedness—namely that the young, educated, high income, and so forth individuals in the representative sample are valid stand-ins for such individuals more generally. The sampling methodology used by Knowledge Networks has been extensively evaluated and shown to meet this assumption (e.g., Yeager et al., 2011), and that is demonstrated in Supplemental Table S1 of the online supplemental material. The second assumption is overlap—that there are individuals at all levels of the moderators in our survey data sets (e.g., that there are at least some very young and very old individuals). This assumption is also shown to be met in Supplemental Table S1 of the supplemental material.

**A behaviorally-informed file drawer adjustment (BIFDA).** Past studies’ effect sizes might have been affected by the so-called “file-drawer problem” (Rosenthal, 1979), which refers to the fact that not all studies that test a given hypothesis end up in the published literature. Even if this is true, the published literature’s average effect sizes will not be distorted if the size and significance of effects in unpublished studies match those in the published studies. However, many investigators presume that file-drawering strengthens effect sizes computed in meta-analyses, because of presumed prejudice against publishing weak or null findings (Munafò et al., 2017). Yet it is also possible that file drawers contain many or even more studies that yielded significant, expected effects, because journals have resisted publishing findings that simply constitute replications of what has already been shown. Therefore, it is impossible to know a priori whether the published literature is misleading, and in what direction a bias might occur.

Common methods for correcting effect sizes for publication bias—such as trim and fill (Duval & Tweedie, 2000), Egger’s test (Egger, Smith, Schneider, & Minder, 1997), or  $p$ -curve (Simonsohn, Nelson, & Simmons, 2014)—do not define a formal model for the probability of selection of studies into the literature. One method that does is the weight-function correction proposed by Vevea (Vevea & Hedges, 1995; Vevea & Woods, 2005; for a discussion of this and other selection models, see Hedges & Vevea, 2005). Researchers can adjust their estimated effect sizes, provided that they know the relative rates at which significant and nonsignificant effects appear in the literature. However, these probabilities are usually not known and are difficult to

estimate precisely in small samples (i.e., fewer than 100 effect sizes; Vevea & Hedges, 1995). To date, this has limited the usefulness of the Vevea method, because the only way to use it is to make assumptions about the magnitude of the bias against publishing null findings or confirmatory replications.

Fortunately, Franco, Malhotra, and Simonovits (2014) recently provided the needed evidence. Franco et al. (2014) began with a “population” of experiments: all of those conducted by the National Science Foundation-funded project called TESS (Time-sharing Experiments for the Social Sciences; www.tessexperiments.org). For more than 15 years, TESS has allowed researchers to conduct experiments with representative samples of American adults, and hundreds of such experiments have been conducted since 2003 by psychologists, sociologists, political scientists, and economists (Franco et al., 2014; TESS, 2018). Interested investigators submitted study proposals to TESS, and a subset of them were approved for implementation, testing new ideas in innovative ways. Once the data were collected by TESS, a public record documented the results of analyses of the data and the ultimate publication status of write-ups of the findings (partly informed by a survey of the designers of the experiments). This permitted Franco et al. (2014) to assess whether findings of significant and nonsignificant effects differed in terms of the rates at which they were published. Only 22% of nonsignificant findings were published in a journal or book, whereas 56% of studies with mixed or strong results were published. (TESS studies were done by scholars from a wide range of disciplines, but publication biases in TESS experiments by psychologists and by political scientists were not significantly different from one another.)

These differences in publication rates can be inputted into the Vevea and Woods (2005) weight-function models to adjust the present study’s meta-analytic effect size estimates.<sup>7</sup> We call this approach the *behaviorally-informed file-drawer adjustment* (BIFDA). We implemented BIFDA for the four effects (conformity, argument quality moderated by need for cognition, false consensus, and ease of retrieval) where conventional significance testing was used and where there were sufficient numbers of studies to implement the Vevea and Woods (2005) corrections.<sup>8,9</sup>

<sup>7</sup> Currently, this can be done using Vevea’s shiny app: <https://vevealab.shinyapps.io/WeightFunctionModel/>.

<sup>8</sup> The probabilistic error studies (the conjunction fallacy and law of large numbers) were not included in the BIFDA because they did not involve significance testing in the original papers, so the  $p < .05$  threshold was not relevant to publication decisions. The lawyer/engineer problem was used too rarely in the past to permit this type of adjustment.

<sup>9</sup> A primary assumption underlying the application of BIFDA is that the rates of censoring effect sizes in the canonical studies is similar to the rates of censoring in studies about which the file drawer can be known. In the present case, this is a conservative assumption. TESS studies were initial tests of new hypotheses, so there is good reason to expect that many such studies will fail to yield statistically significant effects and fail to be deemed publishable. But the phenomena evaluated in our replication studies are classic effects that have been documented repeatedly over decades. Moreover, most of the effects studied here were discovered in an era when collecting and analyzing data was much more costly, burdensome, and slow, and most of the studies that were included in the meta-analyses were conducted in the same era. So  $p$ -hacking and file-drawering seem much less likely. Therefore, basing a BIFDA on the TESS study-failure rate seems likely to yield over-correction, but it is nevertheless a useful, albeit conservative, test of the robustness of our comparisons.

## Results

### Generalizability of the Canonical Studies

**Conformity.** As expected, participants who were told that the majority of Americans favored a government policy ( $M = .321$ ,  $SD = .293$ ) favored the policy more than participants who were told that the majority did not favor the policy ( $M = .295$ ,  $SD = .273$ ),  $F(1, 2119) = 6.66$ ,  $p = .010$ ,  $d = .11$ .

#### Persuasion.

**Argument quality.** As expected, participants were more persuaded by strong arguments than by weak arguments ( $M$  attitude change for strong arguments =  $.072$ ,  $SD = .186$ , and for weak arguments =  $.026$ ,  $SD = .171$ ),  $F(1, 1150) = 17.26$ ,  $p < .001$ ,  $d = .24$ .

**Argument Quality  $\times$  Need for Cognition.** As expected, participants low in need for cognition (at or below the median) were equivalently persuaded by strong and weak arguments against capital punishment ( $M$  attitude change for weak arguments =  $.039$ ,  $SD = .181$ ,  $M$  attitude change for strong arguments =  $.060$ ,  $SD = .190$ ),  $F(1, 586) = 1.35$ ,  $p = .246$ ,  $d = .10$ , whereas participants high in need for cognition (above the median) were more persuaded by strong arguments and were not persuaded by weak ones (weak arguments  $M = .013$ ,  $SD = .158$ , strong arguments  $M = .084$ ,  $SD = .182$ ,  $F(1, 551) = 24.07$ ,  $p < .001$ ,  $d = .43$ ), Argument Quality  $\times$  Need for Cognition interaction,  $F(1, 1148) = 6.75$ ,  $p = .009$ ,  $d = .15$ .<sup>10</sup>

**Base rate underutilization.** As expected, the individuating information manipulation had the expected significant impact on participants' probability judgments (lawyer description  $M = 27.40$ ,  $SD = 29.72$ ; engineer description  $M = 75.66$ ,  $SD = 26.96$ ),  $F(1, 1304) = 951.41$ ,  $p < .001$ ,  $d = 1.71$ , whereas the base rate manipulation had no effect on participants' probability judgments,  $F(1, 1304) = .480$ ,  $p = .480$ ,  $d = .03$ .

**Conjunction fallacy.** As expected, most participants (73%) committed the conjunction fallacy.

**Law of large numbers.** As expected, most participants (72%) did not choose the correct hospital.

**False consensus effect.** As expected, participants who favored a government policy estimated that more people held that opinion than did participants who opposed that government policy (favorers  $M = 47.27\%$ ,  $SD = 16.53$ ; opposers  $M = 39.40\%$ ,  $SD = 15.31$ ,  $F(1, 522) = 25.56$ ,  $p < .001$ ,  $d = .39$ ).

**Ease of retrieval.** As expected, participants found it easier to recall six examples than to recall 12 examples,  $F(1, 1134) = 16.54$ ,  $p < .001$ ,  $d = .24$ . This replicated the manipulation check result in the canonical study (Schwarz et al., 1991). Surprisingly, however, the number of instances recalled did not interact with the valence of those instances when predicting people's ratings of their own assertiveness, Valence  $\times$  Number of Examples interaction  $F(1, 1323) = 1.31$ ,  $p = .251$ ,  $d = .09$ . This is the only instance in which a canonical study's effect was not observed.<sup>11</sup>

### Effect Sizes

**Comparison to past studies.** We located or conducted meta-analyses to produce estimates of the effect sizes for these phenomena based on prior studies with college students (see the online supplemental material for details).<sup>12</sup> The average effect

size in the original study demonstrating each phenomenon was significantly greater than the average effect size produced by meta-analyses of subsequent studies of the same phenomenon conducted with haphazard samples of college students,  $d_s = 0.94$  vs.  $0.66$ ,  $Q(1) = 14.94$ ,  $p < .001$ . To assess heterogeneity across phenomena, we conducted eight separate meta-analytic moderation tests, each one comparing the original effect to the meta-analysis of subsequent studies, and then meta-analyzed those eight moderation tests. This analysis found that the difference was homogeneous across the phenomena studied,  $Q(7) = 1.97$ ,  $p = .96$ .

The average effect size in the meta-analyses of past studies was significantly greater, on average, than was the average effect size in the representative samples,  $d = 0.52$ ,  $Q(1) = 16.41$ ,  $p < .001$  (see Figure 1). A meta-analysis of the moderation tests comparing the meta-analyses to the representative sample effects found no heterogeneity across the phenomena studied,  $Q(7) = 2.05$ ,  $p = .96$ . Sparklines presented in Table 1 show this homogeneity and demonstrate that the result was not driven by the one phenomenon that failed to replicate.

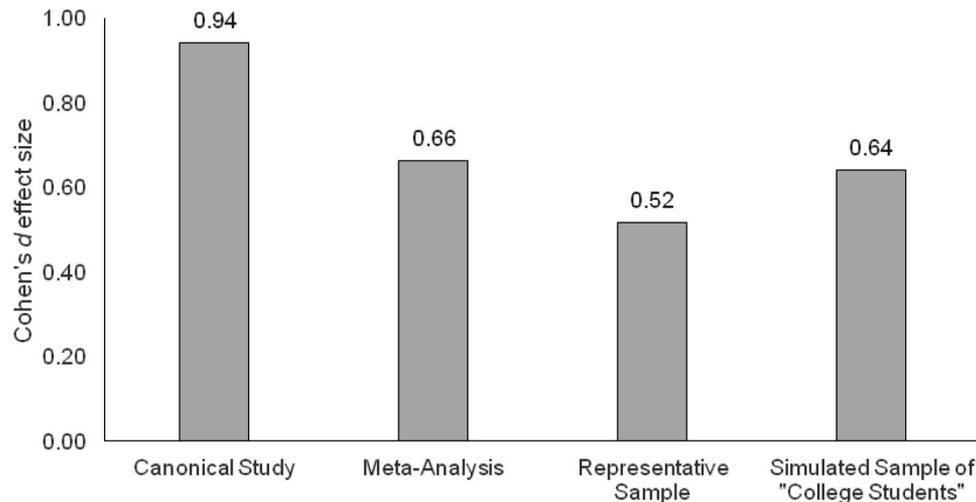
**BIFDA adjustment.** Looking at the four phenomena to which we could apply the BIFDA, the unadjusted meta-analytic average for the college student effect sizes was  $d = .33$ , and the BIFDA-adjusted estimate was  $d = .26$ , a 21% reduction. This suggests that file-drawing strengthened apparent effect sizes in print. Yet the representative sample yielded an effect of  $d = .18$  for these four phenomena, which is still 30% smaller than the BIFDA-adjusted estimate. Thus, it is possible that participants' demographics, not the file drawer, caused discrepant effect sizes.

**Simulating "college student sample" effect sizes.** The statistical simulation estimating the average effect size that would be obtained if the survey participants had the distributions of demographic characteristics of past college student study participants yielded a simulated average effect size of  $d = .64$ . This is significantly larger than the average effect size for the full survey sample,  $Q(1) = 17.63$ ,  $p < .001$ , but not significantly different from the average effect size from meta-analyses of previous studies with haphazard samples of college students,  $Q(1) = 0.717$ ,  $p =$

<sup>10</sup> Although the expected result was obtained when need for cognition was dichotomized at the median (the method used most often in the studies included in Cacioppo, Petty, Feinstein, & Jarvis's (1996) meta-analysis), that result was not observed when coding it differently: When treating need for cognition as a linear continuous variable, when using a GAM to allow non-linear moderation by need for cognition, when dichotomizing need for cognition at the mean, or when trichotomizing it at terciles and dropping the middle tercile. We conducted a  $p$ -curve analysis for the Argument Quality  $\times$  Need for Cognition effects on attitude change (Cacioppo et al., 1983) and found that past studies' results were not likely to be influenced by  $p$ -hacking (studies had evidentiary value,  $Z = -12.38$ ,  $p < .001$ , full disclosure table reported in the online supplemental material). This is consistent with the conclusion that the predicted effect can be observed reliably when splitting need for cognition at the population median.

<sup>11</sup> Exploring the data further, we did not find a significant simple effect of the number of retrieved instances on self-ratings separately within the assertive or unassertive valence conditions (see the online supplemental material).

<sup>12</sup> Our meta-analyses of prior published studies excluded a small number of studies conducted with haphazard samples of noncollege-student adults.



*Figure 1.* Average effect sizes across phenomena tested here. The simulated sample of “college students” refers to estimated effect sizes among people with the distributions of demographic characteristics of students who have been participants in highly cited social psychology studies. The bars represent an unweighted average of all eight studies’ standardized effect sizes, using the statistics reported in Table 1. The effect of the base rate in the lawyer/engineer problem was excluded because that effect was expected to be zero.

.68 (see Figure 1).<sup>13</sup> This suggests that previous college student studies overestimated effect sizes because they included more of the kinds of people who showed those effects more strongly.<sup>14</sup> This result was homogeneous across phenomena, with one exception (see Table 1).

### Moderation by Demographics

Not surprisingly, then, demographics did moderate the magnitudes of the phenomena studied, as shown by meta-analyses of the significance of the effect of each moderator across all phenomena studied (see Table 2 and Figure 2 for aggregate results; see the [online supplemental material](#) for the results of each moderator for each study).<sup>15,16</sup>

**Age.** Age significantly moderated the average effect size across studies,  $Z = 3.64$ ,  $p < .001$ , and this moderation was homogeneous across studies,  $Q(7) = 5.01$ ,  $p = .66$ .<sup>17</sup> Middle-aged adults (25–45) manifested the weakest effects, whereas the youngest (18–25) and oldest adults (45–70) manifested larger effects (see Figure 2). This pattern resembles that identified by Visser and Krosnick (1998) with regard to susceptibility to persuasion.

**Education.** Education also significantly moderated the effects,  $Z = 4.10$ ,  $p < .001$ , and this was homogeneous across studies,  $Q(7) = 3.88$ ,  $p = .79$ . In contrast to claims that more cognitively skilled people are less likely to commit errors of judgment (Kahneman, 2003; Stanovich & West, 1998), such people were more likely to provide heuristic-based responses. Furthermore, more educated participants were significantly more likely to conform and more likely to demonstrate attitude change in response to the stronger arguments (see the [online supplemental material](#)).

**Income.** Income was a significant moderator,  $Z = 3.36$ ,  $p < .001$ , and this was homogeneous across studies,  $Q(7) = 6.77$ ,  $p = .45$ . The highest income group manifested the largest effects across the studies, consistent with theories of the effect of social power on

heuristic responding (Weick & Guinote, 2008) and inconsistent with the theory that higher-power people would demonstrate less attitude change in response to a persuasive manipulation (see Fiske, 2010).

**Sex.** Sex did not moderate the effect size across studies,  $Z = .02$ ,  $p = .98$ , and this nonsignificant effect was homogeneous across studies,  $Q(7) = 5.49$ ,  $p = .60$ . This finding is inconsistent with past studies that have found women to be more likely than men to conform (Cooper, 1979) and that have found women to be more

<sup>13</sup> If a general public sample of adults is more heterogeneous in terms of the outcome variables than a sample of college students, then the variances used to compute the standardized effect sizes in Bar 4 of Figure 1 is bigger than it would be if we had the variance of the hypothetical student participants only. That means that the height of Bar 4 of Figure 1 would be even taller than it is now.

<sup>14</sup> The same conclusions were supported by a simulation done with a slightly different method: generating an estimate of the “upper bound” of the effect size among college students, by estimating the effect that would be expected if an entire sample of participants had the characteristics of the “modal college student.” The simulated average effect size was  $d = .71$ , significantly larger than the average effect size for the full survey sample,  $Q(1) = 39.26$ ,  $p < .001$ , and not significantly different from the average effect size from meta-analyses of previous studies with haphazard samples of college students,  $Q(1) = .51$ ,  $p = .47$ .

<sup>15</sup> We report  $p$  values that are uncorrected for conducting multiple hypothesis tests because there is no consensus in the field on the optimal way to conduct such corrections.

<sup>16</sup> The conformity and persuasion studies were conducted with Sample 1, and the outcome variables were positively correlated,  $r = .26$ ; all other studies were conducted with Sample 2, and the outcomes were very weakly correlated with one another ( $r$  ranged from .01 to .07). Standard errors for all meta-analytic results were corrected for the fact that data came from the same sample using synthetic effect size formulas provided by Borenstein, Hedges, Higgins, and Rothstein (2009).

<sup>17</sup> There are seven degrees of freedom here because the persuasion study contributed two effects.

Table 1  
Effect Sizes for the Present Study and for Studies Conducted Previously

Study	Effects sizes from past studies		The present study's effect sizes		Sparkline Meta-Rep-Sim
	Canonical study	Meta-analysis	Full representative sample	Simulated sample of people with characteristics of college students	
Conformity <sup>a</sup>	.32 [.20, .43]	.25 [.15, .36]	.11 [.03, .19]	.22	
Persuasion <sup>a</sup>					
Argument quality only	1.02 [.54, 1.51]	.77 [.66, .88]	.24 [.13, .36]	.56	
Argument Quality × Need for Cognition	.73 [.32, 1.14]	.29 [.18, .39]	.15 [.04, .27]	.29	
Lawyer/engineer problem <sup>a</sup>					
Individuating information	2.00 [1.63, 2.37]	1.70 [1.42, 1.99]	1.71 [1.53, 1.89]	1.94	
Conjunction fallacy <sup>b</sup>	.88 [.78, .98]	.79 [.78, .82]	.73 [.70, .76]	.76	
Law of large numbers <sup>b</sup>	.78 [.68, .88]	.71 [.68, .73]	.72 [.69, .75]	.68	
False consensus effect <sup>a</sup>	1.37 [.87, 1.87]	.60 [.52, .68]	.39 [.24, .54]	.52	
Ease of retrieval <sup>a</sup>	.42 [.09, .76]	.19 [−.06, .44]	.09 [−.06, .24]	.15	

Note. Numbers in brackets represent 95% confidence intervals.

<sup>a</sup> Effect size = Cohen's  $d$ . <sup>b</sup> Effect size = the proportion of participants providing the incorrect (heuristic) response to the problem. Sparklines visually depict the trends for each phenomenon's effect sizes across the meta-analysis ("Meta"), representative sample ("Rep"), and simulated sample ("Sim"); the sparkline in the header of the table represents the average across all phenomena, so that each study's correspondence to the "V-shaped" trend in the overall result can be gauged.

likely than men to behave and think like lower-power people (see Fiske, 2010).

**Race/ethnicity.** Although race/ethnicity was not a significant moderator of the average effect size across studies,  $Z = .51$ ,  $p = .61$ , this was the only moderator that approached significant heterogeneity across studies (see Table 2), so an exploratory analysis was conducted. In the conformity study, White participants manifested significant conformity,  $t(1874) = 3.12$ ,  $p = .002$ ,  $d = .15$ , whereas a nonsignificant "boomerang" effect of equal size appeared among African American participants,  $t(1874) = -1.17$ ,  $p = .24$ ,  $d = -.16$ . The Poll Result  $\times$  Race interaction was significant,  $F(1, 1874) = 4.49$ ,  $p = .03$ ,  $d = .10$ . In the false consensus effect study, White participants' own attitudes predicted their estimates of how many other Americans had the same attitude,  $t(521) = 5.38$ ,  $p < .001$ ,  $d = .53$ , but non-White participants

showed no false consensus effect,  $t(521) = 0.46$ ,  $p = .65$ ,  $d = .10$ . The Own Attitudes  $\times$  Race interaction was marginally significant,  $F(1, 521) = 3.35$ ,  $p = .07$ ,  $d = .16$ .

Assuming that all participants recognized that the majority of Americans (and therefore the majority of participants in the survey described in the stimulus materials in the conformity experiment) were White, the evidence of conformity by White participants and not among racial or ethnic minority participants is consistent with extant theories that people tend to conform more to others whom they perceive to be similar to them (Hogg, 2010). A similar mechanism may be at work with regard to the false consensus effect: Perhaps people are especially inclined to generalize their own attitudes to groups of other people most similar to them. None of the other studies manifested moderation by race/ethnicity, heterogeneity test,  $Q(5) = 7.33$ ,  $p = .20$ .

**Region.** Effect size was not significantly moderated by region,  $Z = .51$ ,  $p = .60$ , and this result appeared to be homogeneous across studies,  $Q(23) = 24.69$ ,  $p = .37$  (see Table 2). However, because some theory anticipates the most conformity and persuadability among Midwesterners (e.g., Plaut et al., 2002), we explored moderation of each phenomenon by region. As expected, Midwesterners conformed and were persuaded more than people from other regions combined ( $d = .06$  difference),  $Z = 2.81$ ,  $p < .01$ , and these interactions were significantly different from the moderation by region for the other phenomena,  $Q(1) = 7.01$ ,  $p = .01$ . When analyzed individually, none of the other studies manifested significant moderation by region.

## Discussion

Implementing seven classic studies from social psychology in representative samples revealed the following:

1. All but one of the canonical study results appeared in data from large, representative national samples.

Table 2  
Meta-Analyses of Demographic Moderators Across Effects Tested in Representative Samples

Demographic moderator	Effect size ( $r$ )				Heterogeneity		
	for moderator	SE	Z	p	Q	df	p
Sex	.00	.02	-.02	.982	5.49	7	.600
Race	.02	.03	.53	.597	11.58	7	.115
Region	-.01	.01	-.51	.612	24.69	23	.367
Age	.09	.02	3.64	<.001***	5.01	7	.659
Education	.10	.02	4.10	<.001***	3.88	7	.794
Income	.08	.02	3.36	.001***	6.77	7	.454

Note. The meta-analysis of demographic moderation comes from all eight phenomena in Table 1. These tests do not include the effect of the base rate in the lawyer/engineer problem. Standard errors for all meta-analytic results were corrected for those inter-correlations among outcome variables when data come from the same participants, using synthetic effect size formulas provided by Borenstein, Hedges, Higgins, and Rothstein (2009).

\*\*\*  $p < .001$ .

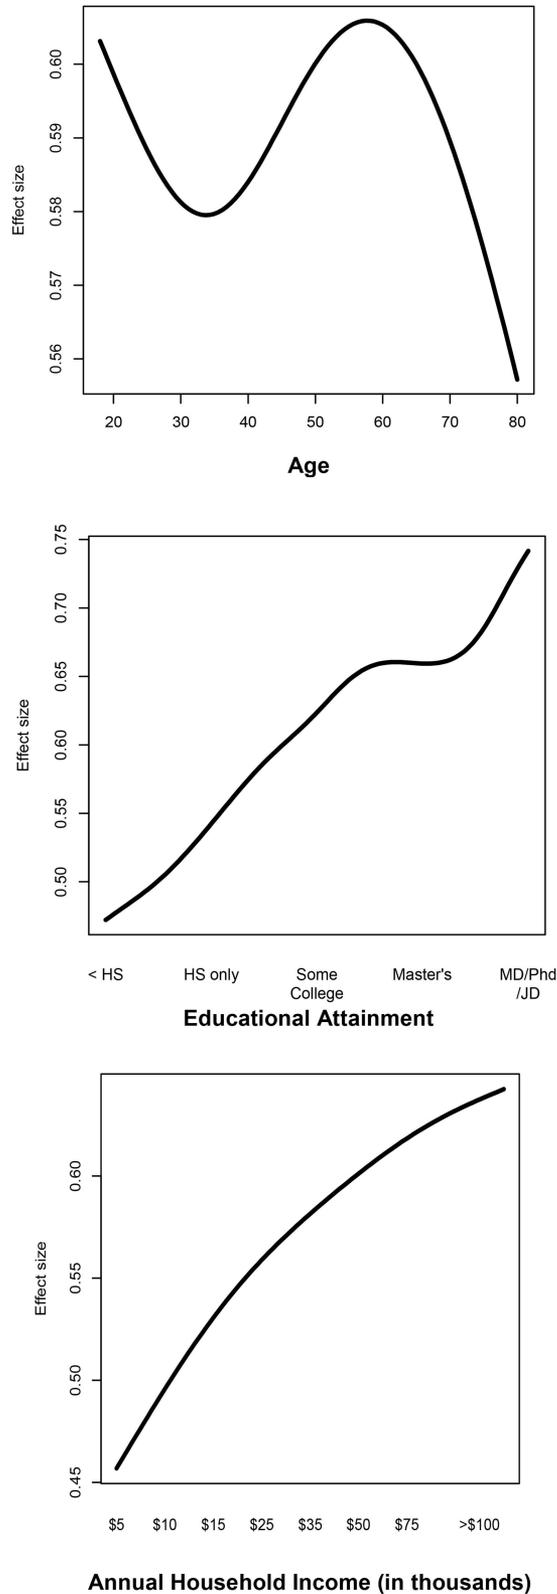


Figure 2. Moderation of effect sizes by age, education, and income for the phenomena tested here.

2. Most demographics moderated the effect sizes. The directions of the moderation were sometimes in line with expectations based on theory or prior research and sometimes not.
3. The largest effects were generally found among people with the characteristics of the college students who have usually partaken in psychological experiments in labs (young, wealthy, and well-educated).
4. Effect sizes in the canonical studies were larger than meta-analytic effects sizes documented in subsequent studies of college students, and the latter effect sizes were larger than the effect sizes in representative samples of American adults.
5. The second of these two differences was not fully attributable to file-drawing of nonsignificant study results, as shown by BIFDA.
6. Simulated effect sizes among Americans with the characteristics of college student experiment participants closely resembled the meta-analytic effect sizes from previous studies of college student participants.

The reproduction of canonical effects constitutes important evidence that foundational phenomena observed in past studies of college students also occur in the general public as a whole. It was not a foregone conclusion that these effects would appear here: We are an independent research team that used different types of participants, sometimes different stimuli, different data collection methods, no file-drawing, and an approach designed to minimize researcher degrees of freedom. Our replications are reassuring about the internal and external validity of past studies and suggest that past studies might not have been as afflicted by questionable research practices as some observers have claimed.

Like us, Mullinix, Leeper, Druckman, and Freese (2015) also found that three framing experiments with college student samples were replicated in national representative samples of adults. However, Mullinix et al. (2015) reached a different conclusion than we did regarding effect sizes; they found no differences across the types of samples, whereas we did. Because the essence of these framing effects is opinion change, and because we found more opinion change among college students and among adults with the characteristics of college students than among a general public sample, we would expect to see the same pattern in Mullinix et al.'s (2015) data. However, Mullinix et al. (2015) studied new experimental manipulations developed for their investigation, whereas we examined classic phenomena documented in numerous previous studies. Nonetheless, we see no reason why this should moderate the difference in effect sizes between samples. We therefore look forward to future research exploring this discrepancy in findings.

### Relation to Other Replication Efforts

The present high rate of replication (86%) aligns with research by Klein et al. (2014), who found a replication rate of 85% when conducting a hand-picked group of classic and contemporary experiments with haphazard samples of undergraduates and volun-

teers on Amazon's Mechanical Turk. Thus, our findings are, in this sense, in line with Klein et al.'s (2014) speculation that "replicability is more dependent on the effect itself than on the sample and setting used to investigate the effect" (p. 142). However, the present finding that effect sizes varied by sample composition suggests that perhaps that speculation is better limited to whether an effect is observed rather than to its strength.

Our high replication rate might seem to contradict evidence that only about one third of 100 recent social and cognitive psychology findings could be replicated (Open Science Collaboration, 2015).<sup>18</sup> However, those replication attempts focused on new phenomena that had not yet been subjected to extensive efforts to replicate them. In contrast, the present article focuses on effects that have been replicated many times with college student participants and therefore had a high likelihood a priori of being observed again. Furthermore, due to the large sample sizes used in the studies reported here, the resulting statistical tests had power that approached 1.0, whereas more than 20% of the Open Science Collaboration's replications used smaller samples than the original studies had (also see Gilbert, King, Pettigrew, & Wilson, 2016) and were conducted by undergraduate or graduate researchers with various levels of expertise. Therefore, it may be best to view the Open Science Collaboration (2015) replication rate as describing new discoveries published for the first time with modestly powered investigations and many different researchers, whereas the present findings regard canonical phenomena tested in high-powered studies, with generalizable samples and data collected by a single professional research firm.

The findings in Figure 1 resonate with the notion that effect sizes get weaker over time, also called the "discoverer's curse" or the "decline effect" (Bakker, van Dijk, & Wicherts, 2012; Ioannidis, 2005; Jennions & Møller, 2002; Munafò et al., 2017; Schooler, 2011). This may occur because of the use of data collection strategies and/or analytic techniques that misleadingly enhance apparent effect sizes, including publication bias, studying small samples, selective reporting, *p*-hacking, accidental errors in analysis, and more (Bakker & Wicherts, 2011; Francis, 2012; Franco et al., 2014; Hartgerink, van Aert, Nuijten, Wicherts, & van Assen, 2016; Ioannidis, 2005; John, Loewenstein, & Prelec, 2012; Schimmack, 2012). Consistent with the notion of a decline effect, we observed stronger effects in the most cited studies of phenomena than in subsequent studies of them.

Going beyond past research showing decline effects, we showed a second decline effect, when studies are replicated in representative samples rather than haphazard samples. In that regard, the present evidence resonates with findings reported by Trzesniewski and Donnellan (2010), who demonstrated that the conclusions of meta-analyses of data from haphazard samples of college students do not routinely accurately describe the strengths of relations between variables in the population.

### Explaining Differences in Effect Sizes

What caused the decline in effect size from studies of college students to the national sample (see Figure 1)? One possibility, empirically supported in this case, is the difference between the people who participated in the studies. Using the survey data to simulate effect sizes among participants with the characteristics of typical college student participants in lab studies yielded effect

sizes that were stronger than effects in the full representative sample. This comparison has strong internal validity, because all aspects of methodology were held constant except the sample. Furthermore, the simulation yielded effect sizes that were comparable to those found in meta-analyses of lab studies (i.e., comparing the second bar to the fourth bar in Figure 1).

Other explanations for the difference in effect sizes between Bars 2 and 3 in Figure 1 are possible as well, but the data suggest these explanations are not likely in the present case. For example, the meta-analytic effect sizes characterized by Bar 2 might have been inflated due to questionable research practices (John et al., 2012). Reassuring in this regard is the fact that *p*-curve analyses of past literature did not suggest reasons for concern (disclosure tables presented in the online supplemental material) although such *p*-curve analysis results alone cannot rule out this possibility (Hartgerink et al., 2016; Simonsohn, Simmons, & Nelson, 2015).

Smaller effect sizes in national samples might simply be a statistical artifact, because of more variance in the outcome variables in those samples caused by greater heterogeneity of participant characteristics, which alone would reduce standardized effect sizes. But the present studies that allowed for comparison of unstandardized or raw metric effect sizes (law of large numbers, conjunction fallacy, lawyer/engineer, and false consensus) revealed the same results as did the standardized measures: the national sample effects were smaller on average than the effects observed in past studies. For instance, the false consensus effect was an 8 percentage-point difference here and 17 points in Ross et al.'s (1977) study.

Another possible explanation involves calibration of manipulations (Gilbert et al., 2016; Schooler, 2014; Wilson, Aronson, & Carlsmith, 2010). For example, in the eyes of college-student-like national sample members, the strong and weak arguments used in the present persuasion experiment might seem quite different in terms of quality, but to middle-age adults, the arguments might not appear to be as different in terms of quality. In keeping with methodological recommendations with respect to measurement invariance made by Vandenberg and Lance (2000), we conducted follow-up experiments with large nationally representative samples to test whether the materials used in the current study varied in their evocativeness or interpretations in ways that might explain the demographic moderation of the primary studies observed here. In no case was the obtained evidence consistent with that claim (see the online supplemental material). So the demographic differences documented here seem unlikely to be attributable to differences across demographic groups in the effectiveness or calibration of the manipulations.

### The Ease of Retrieval Effect

The most recently discovered canonical finding failed to appear in the large representative sample—the ease of retrieval effect (Schwarz et al., 1991). This phenomenon has been explored in many past studies (Greifeneder, Bless, & Pham, 2011; Weingarten & Hutchinson, 2018), but almost none of those studies tested for

<sup>18</sup> Replication rates for the Open Science Collaboration (2015) were much higher when using other criteria, such as whether the replication effect size fell within the original effect size's 95% confidence interval, or whether the two effects were meta-analytically significant.

the theoretically critical but marginally significant cross-over interaction reported in the canonical publication, and the few instances that did seek to replicate that interaction failed to do so with college student participants (von Helversen, Gendolla, Winkielman, & Schmidt, 2008; Vaughn, 1998). Therefore, our failure to observe that interaction in the national data might best be viewed as successful generalization of these later failures to replicate.

### Alternative Social-Psychological Phenomena

The present evidence does not support the conclusion that all social psychological phenomena are stronger among college students than in representative samples. To reach such a conclusion, we would have had to define a “population” of social psychological phenomena and randomly select a large set of such phenomena to test. So the present findings should be considered a first step in a programmatic effort to explore the generalizability of social psychological phenomena to well-defined populations of people who are sampled truly randomly for investigation.

The large “population” of documented social psychological phenomena is the result of researchers’ choices of what to study and what to publish, and these choices have presumably been influenced in part by how easy it has been to document an effect among the most frequently studied participants (i.e., college students). If instead, the field had followed the pathway of some subfields of political science and sociology and economics by studying random samples of adult populations, we might have ended up with a different set of phenomena dominating our literature, and those phenomena might be stronger among the general adult population (with whom the phenomena were originally identified) than among college students.

For example, working-class individuals may hold different cultural values than middle- or upper-class individuals, and these differing values might alter motivations for and styles of thinking and reasoning (Kraus, Piff, Mendoza-Denton, Rheinschmidt, & Keltner, 2012; Stephens, Markus, & Phillips, 2014). Thus, the social-cognition literature might look very different if classic studies had been conducted initially in working class communities.

### More Use of Representative Samples

Our findings might be viewed as discouraging researchers from routinely testing the generalizability of findings from haphazard samples of students or adults to representative, random samples of adults. After all, we found generalization of all of the examined phenomena (real effects and null ease-of-retrieval interaction). But if researchers wish to take seriously the effect sizes documented with haphazard samples, our findings suggest that occasional tests of phenomena in representative samples will have scientific and practical value.

Haphazard samples of adults are unlikely to be sufficient for testing generalizability. Such samples differ from one another in uncontrolled ways in terms of demographics. Because such demographics often moderated effect sizes in the present investigation, uncontrolled variation across studies can create the illusion of failure to replicate when the differing results are actually due to systematic differences in sample composition. As Paolacci and Chandler (2014) concluded, M-Turk study participants “should not

be treated as representative of the general population” (p. 185), because they manifest considerable systematic biases similar to the characteristics of college student study participants: younger, more educated, underrepresenting African Americans and Hispanics, more liberal, and more. Imposing quotas on demographics or weighting nonprobability samples using demographics does not solve this problem (MacInnis et al., 2018; Yeager et al., 2011).

### The Value of Demographic Moderation for Theory Development

The present evidence of demographic moderation offers an opportunity to advance theory development, in part because we failed to observe evidence consistent with some past speculations about demographic moderators. Consider, for example, Kahneman’s (2003) speculation that heuristics and biases in reasoning might be most often apparent when “System 2” (the slow, deliberate, thoughtful system) is compromised, as it might be among individuals with the most limited cognitive skills. Previous research leading to that speculation came from examining variance in SAT scores among samples of undergraduates attending highly regarded colleges (e.g., Stanovich & West, 1998). But in such settings, there is a restriction of range as compared to the entire population and, because the SAT is a criterion for admission to educational institutions, the joint distribution between SAT scores and the underlying psychological mechanisms is likely to be distorted by “collider bias” (Morgan & Winship, 2014).

The present study’s use of representative samples avoids the bias caused by examining a selective subset of people in a restricted age range. The present evidence runs opposite to Kahneman’s speculation. The largest effects of the heuristic manipulations were observed among the most educated participants, who had the most advanced cognitive skills (Brinch & Galloway, 2012; Ceci, 1991; Rietveld et al., 2014). Making matters worse, the Open Science Collaboration (2015) failed to replicate the predicted moderation of heuristics and biases by cognitive skills, so this issue clearly deserves more investigation.

These findings and others like them may inspire reconsideration of core elements of long-standing theories and may also prove practically relevant. Insights from social psychology are increasingly being repackaged as behavioral economics and then incorporated into high stakes decisions by companies and governments. For instance, conformity to a descriptive norm, a finding replicated here, is a core tenet of Opower™, a behavioral energy management company (Allcott, 2015). Yet descriptive norm information had a “boomerang” effect among African Americans that was equal in magnitude to the positive and significant effect among White participants in the present study, perhaps because African Americans are construing the norm as applying to out-group members who may harbor negative stereotypes about their group. Thus, it is interesting to note that social groups that have traditionally been marginalized by societies generally may have also been unintentionally marginalized in social psychology because of a tendency to presume generalizability without testing it. Our evidence suggests value in more such testing of moderation by demographics in the spirit of social justice.

## Conclusion

This is a feel-good article for social psychologists, who like replication, generalization, and theoretical advancement. This article provides all three. The fact that almost all effects appeared in representative samples of American adults documents the replicability of the effects and their generalizability to the general population, just as most psychologists have assumed for decades.

The evidence of moderation is among the first empirical testing called for by Sears (1986) in his landmark article decades ago and suggests the possibility of fruitful bridging of the agenda of social psychologists with that of sociologists, who have a special interest in demographics and much rich theory on their origins and effects. Future study of the foundational effects studied here through such a sociological lens may lead to interesting and useful insights into the functioning of these and many other social psychological processes.

This article's agenda to explore generalization of effects across population subgroups complements other work done to date exploring generalization of social psychological phenomena across cultures (e.g., Henrich et al., 2010) and over time (e.g., Twenge, Konrath, Foster, Campbell, & Bushman, 2008) and encourages that such assessments be done rigorously in the future using representative samples rather than haphazard ones (e.g., De Neve et al., 2018; Trzesniewski & Donnellan, 2010). Such a research agenda need not be cost-prohibitive. Random sampling of the American adult population has been used in studies appearing in psychology journals thanks to TESS, which provided the data necessary for the BIFDA analysis above. Thus, the sort of investigation reported in this article can easily be conducted by investigators in the future, at no cost to them, by testing their effects on the TESS platform.

In sum, we hope that the present findings encourage investigators to occasionally conduct tests of classic and novel effects in random samples. This may help the field interrogate the heterogeneity of those effects across groups of individuals and deepen our understanding of their mechanisms.

## References

- Allcott, H. (2015). Site selection bias in program evaluation. *The Quarterly Journal of Economics*, *130*, 1117–1165. <http://dx.doi.org/10.1093/qje/qjv015>
- Andersen, R. (2009). Nonparametric methods for modeling nonlinearity in regression analysis. *Annual Review of Sociology*, *35*, 67–85. <http://dx.doi.org/10.1146/annurev.soc.34.040507.134631>
- Asch, S. (1952). *Social psychology*. New York, NY: Prentice Hall. <http://dx.doi.org/10.1037/10025-000>
- Bakker, M., van Dijk, A., & Wicherts, J. M. (2012). The rules of the game called psychological science. *Perspectives on Psychological Science*, *7*, 543–554. <http://dx.doi.org/10.1177/1745691612459060>
- Bakker, M., & Wicherts, J. M. (2011). The (mis)reporting of statistical results in psychology journals. *Behavior Research Methods*, *43*, 666–678. <http://dx.doi.org/10.3758/s13428-011-0089-5>
- Berinsky, A. J., Huber, G. A., & Lenz, G. S. (2012). Evaluating online labor markets for experimental research: Amazon.com's Mechanical Turk. *Political Analysis*, *20*, 351–368. <http://dx.doi.org/10.1093/pan/mpm057>
- Bond, R., & Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using Asch's (1952b, 1956) line judgment task. *Psychological Bulletin*, *119*, 111–137. <http://dx.doi.org/10.1037/0033-2909.119.1.111>
- Borenstein, M., Hedges, L. V., Higgins, J. P. T., & Rothstein, H. R. (2009). *Introduction to meta-analysis*. Hoboken, NJ: Wiley. <http://dx.doi.org/10.1002/9780470743386.ch20>
- Brinch, C. N., & Galloway, T. A. (2012). Schooling in adolescence raises IQ scores. *Proceedings of the National Academy of Sciences of the United States of America*, *109*, 425–430. <http://dx.doi.org/10.1073/pnas.1106077109>
- Buhrmester, M., Kwang, T., & Gosling, S. D. (2011). Amazon's Mechanical Turk: A new source of inexpensive, yet high-quality, data? *Perspectives on Psychological Science*, *6*, 3–5. <http://dx.doi.org/10.1177/1745691610393980>
- Bureau of Labor Statistics, U.S. Department of Labor. (2014). *National Longitudinal Survey of Youth 1979 cohort, 1979–2012 (rounds 1–25)*. Retrieved from <https://www.bls.gov/nls/nlsy79.htm>
- Cacioppo, J. T., Petty, R. E., Feinstein, J. A., & Jarvis, W. B. G. (1996). Dispositional differences in cognitive motivation: The life and times of individuals varying in need for cognition. *Psychological Bulletin*, *119*, 197–253. <http://dx.doi.org/10.1037/0033-2909.119.2.197>
- Cacioppo, J. T., Petty, R. E., Kao, C. F., & Rodriguez, R. (1986). Central and peripheral routes to persuasion: An individual difference perspective. *Journal of Personality and Social Psychology*, *51*, 1032–1043. <http://dx.doi.org/10.1037/0022-3514.51.5.1032>
- Cacioppo, J. T., Petty, R. E., & Morris, K. J. (1983). Effects of need for cognition on message evaluation, recall, and persuasion. *Journal of Personality and Social Psychology*, *45*, 805–818. <http://dx.doi.org/10.1037/0022-3514.45.4.805>
- Ceci, S. J. (1991). How much does schooling influence general intelligence and its cognitive components? A reassessment of the evidence. *Developmental Psychology*, *27*, 703–722. <http://dx.doi.org/10.1037/0012-1649.27.5.703>
- Cialdini, R. B. (2003). Crafting normative messages to protect the environment. *Current Directions in Psychological Science*, *12*, 105–109. <http://dx.doi.org/10.1111/1467-8721.01242>
- Cook, T. D., & Groom, C. (2004). The methodological assumptions of social psychology: The mutual dependence of substantive theory and method choice. In C. Sansone, C. C. Morf, & A. T. Panter (Eds.), *The Sage handbook of methods in social psychology* (pp. 19–44). Thousand Oaks, CA: Sage.
- Cooper, H. M. (1979). Statistically combining independent studies: A meta-analysis of sex differences in conformity research. *Journal of Personality and Social Psychology*, *37*(1), 131–146. <http://dx.doi.org/10.1037/0022-3514.37.1.131>
- Crump, M. J. C., McDonnell, J. V., & Gureckis, T. M. (2013). Evaluating Amazon's Mechanical Turk as a tool for experimental behavioral research. *PLoS ONE*, *8*(3), e57410. <http://dx.doi.org/10.1371/journal.pone.0057410>
- De Neve, J.-E., Ward, G., De Keulenaer, F., Van Landeghem, B., Kavetsos, G., & Norton, M. I. (2018). The asymmetric experience of positive and negative economic growth: Global evidence using subjective well-being data. *The Review of Economics and Statistics*, *100*, 362–375. [http://dx.doi.org/10.1162/REST\\_a\\_00697](http://dx.doi.org/10.1162/REST_a_00697)
- Duval, S., & Tweedie, R. (2000). A nonparametric “trim and fill” method of accounting for publication bias in meta-analysis. *Journal of the American Statistical Association*, *95*, 89–98.
- Eagly, A. H., & Warren, R. (1976). Intelligence, comprehension, and opinion change. *Journal of Personality*, *44*, 226–242. <http://dx.doi.org/10.1111/j.1467-6494.1976.tb00120.x>
- Eaton, A. A., Visser, P. S., Krosnick, J. A., & Anand, S. (2009). Social power and attitude strength over the life course. *Personality and Social Psychology Bulletin*, *35*, 1646–1660. <http://dx.doi.org/10.1177/0146167209349114>
- Edwards-Levy, A. (2017, May 5). *What went wrong with last year's election surveys? Pollsters have some answers*. Retrieved from <https://>

- [www.huffingtonpost.com/entry/polls-wrong-2016\\_us\\_590b9e9de4b0104c734d6132](http://www.huffingtonpost.com/entry/polls-wrong-2016_us_590b9e9de4b0104c734d6132)
- Egger, M., Smith, G. D., Schneider, M., & Minder, C. (1997). Bias in meta-analysis detected by a simple, graphical test. *British Medical Journal (Clinical Research Ed.)*, *315*, 629–634. <http://dx.doi.org/10.1136/bmj.315.7109.629>
- Erikson, E. (1968). *Identity: Youth and crisis*. New York, NY: Norton & Co.
- Feller, A., & Holmes, C. C. (2009). *Beyond topline: Heterogeneous treatment effects in randomized experiments*. Unpublished Manuscript, Department of Statistics, University of Oxford, England.
- Fiske, S. T. (2010). Interpersonal stratification: Status, power, and subordination. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (5th ed., pp. 941–982). Hoboken, NJ: Wiley. <http://dx.doi.org/10.1002/9780470561119.socpsy002026>
- Francis, G. (2012). Too good to be true: Publication bias in two prominent studies from experimental psychology. *Psychonomic Bulletin & Review*, *19*, 151–156. <http://dx.doi.org/10.3758/s13423-012-0227-9>
- Franco, A., Malhotra, N., & Simonovits, G. (2014). Publication bias in the social sciences: Unlocking the file drawer. *Science*, *345*, 1502–1505. <http://dx.doi.org/10.1126/science.1255484>
- Gilbert, D. T., King, G., Pettigrew, S., & Wilson, T. D. (2016). Comment on “Estimating the reproducibility of psychological science.” *Science*, *351*, 1037. <http://dx.doi.org/10.1126/science.aad7243>
- Goodman, J. K., Cryder, C. E., & Cheema, A. (2013). Data collection in a flat world: The strengths and weaknesses of Mechanical Turk samples. *Journal of Behavioral Decision Making*, *26*, 213–224. <http://dx.doi.org/10.1002/bdm.1753>
- Gosling, S. D., Vazire, S., Srivastava, S., & John, O. P. (2004). Should we trust web-based studies? A comparative analysis of six preconceptions about internet questionnaires. *American Psychologist*, *59*, 93–104. <http://dx.doi.org/10.1037/0003-066X.59.2.93>
- Green, D. P., & Kern, H. L. (2012). Modeling heterogeneous treatment effects in survey experiments with Bayesian additive regression trees. *Public Opinion Quarterly*, *76*, 491–511. <http://dx.doi.org/10.1093/poq/nfs036>
- Greifeneder, R., Bless, H., & Pham, M. T. (2011). When do people rely on affective and cognitive feelings in judgment? A review. *Personality and Social Psychology Review*, *15*, 107–141. <http://dx.doi.org/10.1177/1088868310367640>
- Hanretty, C. (2016, June 24). *Here's why pollsters and pundits got Brexit wrong*. Retrieved from [https://www.washingtonpost.com/news/monkey-cage/wp/2016/06/24/heres-why-pollsters-and-pundits-got-brexit-wrong/?noredirect=on&utm\\_term=.568e1cdf8403](https://www.washingtonpost.com/news/monkey-cage/wp/2016/06/24/heres-why-pollsters-and-pundits-got-brexit-wrong/?noredirect=on&utm_term=.568e1cdf8403)
- Hartgerink, C. H. J., van Aert, R. C. M., Nuijten, M. B., Wicherts, J. M., & van Assen, M. A. L. M. (2016). Distributions of p-values smaller than .05 in psychology: What is going on? *PeerJ*, *4*, e1935. <http://dx.doi.org/10.7717/peerj.1935>
- Hedges, L. V., & Vevea, J. (2005). Selection method approaches. In H. R. Rothstein, A. J. Sutton, & M. Borenstein (Eds.), *Publication bias in meta-analysis: Prevention, assessment, and adjustments* (pp. 145–174). Chichester, UK: Wiley, Ltd.
- Henrich, J., Heine, S. J., & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences*, *33*(2–3), 61–83. <http://dx.doi.org/10.1017/S0140525X0999152X>
- Henry, P. J. (2008). College sophomores in the laboratory redux: Influences of a narrow data base on social psychology's view of the nature of prejudice. *Psychological Inquiry*, *19*, 49–71. <http://dx.doi.org/10.1080/10478400802049936>
- Hertwig, R., & Chase, V. M. (1998). Many reasons or just one: How response mode affects reasoning in the conjunction problem. *Thinking & Reasoning*, *4*, 319–352. <http://dx.doi.org/10.1080/135467898394102>
- Hogg, M. A. (2010). Influence and leadership. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (5th ed., Vol. 2, pp. 1167–1207). Hoboken, NJ: Wiley. <http://dx.doi.org/10.1002/9780470561119.socpsy002031>
- Ioannidis, J. P. A. (2005). Why most published research findings are false. *PLoS ONE*, *18*, e124. <http://dx.doi.org/10.1371/journal.pmed.0020124>
- Jennions, M. D., & Møller, A. P. (2002). Relationships fade with time: A meta-analysis of temporal trends in publication in ecology and evolution. *Proceedings of the Royal Society of London, Series B: Biological Sciences*, *269*, 43–48. <http://dx.doi.org/10.1098/rspb.2001.1832>
- John, L. K., Loewenstein, G., & Prelec, D. (2012). Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological Science*, *23*, 524–532. <http://dx.doi.org/10.1177/0956797611430953>
- Kahneman, D. (2003). Maps of bounded rationality: Psychology for behavioral economics. *The American Economic Review*, *93*, 1449–1475. <http://dx.doi.org/10.1257/000282803322655392>
- Kahneman, D. (2011). *Thinking, fast and slow*. New York, NY: Farrar, Straus & Giroux.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, *80*, 237–251. <http://dx.doi.org/10.1037/h0034747>
- Keele, L. J. (2008). *Semiparametric regression for the social sciences*. Hoboken, NJ: Wiley.
- Keeter, S., Kennedy, C., Dimock, M., Best, J., & Craighill, P. (2006). Gauging the impact of growing nonresponse on estimates from a national RDD telephone survey. *Public Opinion Quarterly*, *70*, 759–779. <http://dx.doi.org/10.1093/poq/nfl035>
- King, L., & Sobelman, B. (2015, March 18). *How did the polls in Israel get it so wrong?* Retrieved from <http://www.latimes.com/world/middleeast/la-fg-israel-polls-wrong-20150318-story.html>
- 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>
- Kraus, M. W., Piff, P. K., Mendoza-Denton, R., Rheinschmidt, M. L., & Keltner, D. (2012). Social class, solipsism, and contextualism: How the rich are different from the poor. *Psychological Review*, *119*, 546–572. <http://dx.doi.org/10.1037/a0028756>
- MacInnis, B., Krosnick, J. A., Ho, A. S., & Cho, M.-J. (2018). The accuracy of measurements with probability and nonprobability survey samples: Replication and extension. *Public Opinion Quarterly*, *82*, 707–744. <http://dx.doi.org/10.1093/poq/nfy038>
- McLean, A., Krosnick, J. A., & Tahk, A. (2018). *Accuracy of national and state polls in predicting the outcome of the 2016 presidential election*. Unpublished manuscript, Stanford University, Stanford, CA.
- Morgan, S. L., & Winship, C. (2014). *Counterfactuals and causal inference*. New York, NY: Cambridge University Press. <http://dx.doi.org/10.1017/CBO9781107587991>
- Mullen, B., Atkins, J. L., Champion, D. S., Edwards, C., Hardy, D., Story, J. E., & Vanderklok, M. (1985). The false consensus effect: A meta-analysis of 115 hypothesis tests. *Journal of Experimental Social Psychology*, *21*, 262–283. [http://dx.doi.org/10.1016/0022-1031\(85\)90020-4](http://dx.doi.org/10.1016/0022-1031(85)90020-4)
- Mullinix, K. J., Leeper, T. J., Druckman, J. N., & Freese, J. (2015). The generalizability of survey experiments. *Journal of Experimental Political Science*, *2*, 109–138. <http://dx.doi.org/10.1017/XPS.2015.19>
- Munafò, M. R., Nosek, B. A., Bishop, D. V. M., Button, K. S., Chambers, C. D., Percie du Sert, N., . . . Ioannidis, J. P. A. (2017). A manifesto for reproducible science. *Nature Human Behaviour*, *1*, 0021. <http://dx.doi.org/10.1038/s41562-016-0021>
- Newkirk, V. R., II. (2016, November 9). *What went wrong with the 2016 poll?* Retrieved from <https://www.theatlantic.com/politics/archive/2016/11/what-went-wrong-polling-clinton-trump/507188/>
- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, *349*(6251), aac4716. <http://dx.doi.org/10.1126/science.aac4716>

- Paolacci, G., & Chandler, J. (2014). Inside the Turk: Understanding Mechanical Turk as a participant pool. *Current Directions in Psychological Science*, 23, 184–188. <http://dx.doi.org/10.1177/0963721414531598>
- Peterson, R. A. (2001). On the use of college students in social science research: Insights from a second-order meta-analysis. *Journal of Consumer Research*, 28, 450–461. <http://dx.doi.org/10.1086/323732>
- Petty, R. E., & Cacioppo, J. T. (1996). Addressing disturbing and disturbed consumer behavior: Is it necessary to change the way we conduct behavioral science? *Journal of Marketing Research*, 33, 1–8. <http://dx.doi.org/10.1177/002224379603300101>
- Plaut, V. C., Markus, H. R., & Lachman, M. E. (2002). Place matters: Consensual features and regional variation in American well-being and self. *Journal of Personality and Social Psychology*, 83, 160–184. <http://dx.doi.org/10.1037/0022-3514.83.1.160>
- Reis, H. T., & Gosling, S. D. (2010). Social psychological methods outside the laboratory. In S. T. Fiske, D. T. Gilbert, & G. Lindzey (Eds.), *Handbook of social psychology* (5th ed., pp. 82–114). Hoboken, NJ: Wiley. <http://dx.doi.org/10.1002/9780470561119.socpsy001003>
- Rietveld, C. A., Esko, T., Davies, G., Pers, T. H., Turley, P., Benyamin, B., . . . Koellinger, P. D. (2014). Common genetic variants associated with cognitive performance identified using the proxy-phenotype method. *Proceedings of the National Academy of Sciences of the United States of America*, 111, 13790–13794. <http://dx.doi.org/10.1073/pnas.1404623111>
- Rosenthal, R. (1979). The file drawer problem and tolerance for null results. *Psychological Bulletin*, 86, 638–641. <http://dx.doi.org/10.1037/0033-2909.86.3.638>
- Ross, L., Greene, D., & House, P. (1977). The “false consensus effect”: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology*, 13, 279–301. [http://dx.doi.org/10.1016/0022-1031\(77\)90049-X](http://dx.doi.org/10.1016/0022-1031(77)90049-X)
- 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>
- Schooler, J. W. (2011). Unpublished results hide the decline effect. *Nature*, 470, 437. <http://dx.doi.org/10.1038/470437a>
- Schooler, J. W. (2014). Turning the lens of science on itself: Verbal overshadowing, replication, and metascience. *Perspectives on Psychological Science*, 9, 579–584. <http://dx.doi.org/10.1177/1745691614547878>
- Schwarz, N., Bless, H., Strack, F., Klumpp, G., Rittenauer-Schatka, H., & Simons, A. (1991). Ease of retrieval as information: Another look at the availability heuristic. *Journal of Personality and Social Psychology*, 61, 195–202. <http://dx.doi.org/10.1037/0022-3514.61.2.195>
- Sears, D. O. (1986). College sophomores in the laboratory: Influences of a narrow data base on social psychology's view of human nature. *Journal of Personality and Social Psychology*, 51, 515–530. <http://dx.doi.org/10.1037/0022-3514.51.3.515>
- Sherif, M. (1936). *The psychology of social norms*. New York, NY: Harper & Brothers.
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science*, 22, 1359–1366. <http://dx.doi.org/10.1177/0956797611417632>
- 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. P., & Nelson, L. D. (2015). Better P-curves: Making P-curve analysis more robust to errors, fraud, and ambitious P-hacking, a Reply to Ulrich and Miller (2015). *Journal of Experimental Psychology: General*, 144, 1146–1152. <http://dx.doi.org/10.1037/xge0000104>
- Stanovich, K. E., & West, R. F. (1998). Individual differences in rational thought. *Journal of Experimental Psychology: General*, 127, 161–188. <http://dx.doi.org/10.1037/0096-3445.127.2.161>
- Stephens, N. M., Markus, H. R., & Phillips, L. T. (2014). Social class culture cycles: How three gateway contexts shape selves and fuel inequality. *Annual Review of Psychology*, 65, 611–634. <http://dx.doi.org/10.1146/annurev-psych-010213-115143>
- Time-sharing Experiments for the Social Sciences (TESS). (2018, August 26). *Information about some previous TESS studies*. Retrieved from <http://www.tessexperiments.org/previousstudies.html>
- Trzesniewski, K. H., & Donnellan, M. B. (2010). Rethinking “generation me”: A study of cohort effects from 1976–2006. *Perspectives on Psychological Science*, 5, 58–75. <http://dx.doi.org/10.1177/1745691609356789>
- 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>
- Tversky, A., & Kahneman, D. (1983). Extensional versus intuitive reasoning: The conjunction fallacy in probability judgment. *Psychological Review*, 90, 293–315. <http://dx.doi.org/10.1037/0033-295X.90.4.293>
- Twenge, J. M., Konrath, S., Foster, J. D., Campbell, W. K., & Bushman, B. J. (2008). Egos inflating over time: A cross-temporal meta-analysis of the Narcissistic Personality Inventory. *Journal of Personality*, 76, 875–902. <http://dx.doi.org/10.1111/j.1467-6494.2008.00507.x>
- Vandenberg, R. J., & Lance, C. E. (2000). A review and synthesis of the measurement invariance literature: Suggestions, practices, and recommendations for organizational research. *Organizational Research Methods*, 3, 4–70. <http://dx.doi.org/10.1177/109442810031002>
- Van Lange, P. A. M., Schippers, M., & Balliet, D. (2011). Who volunteers in psychology experiments? An empirical review of prosocial motivation in volunteering. *Personality and Individual Differences*, 51, 279–284. <http://dx.doi.org/10.1016/j.paid.2010.05.038>
- Vaughn, L. A. (1998). *Expertise and use of experienced ease or difficulty of recall for social judgments* (Doctoral dissertation). University of Michigan, Ann Arbor, MI.
- Vevea, J. L., & Hedges, L. V. (1995). A general linear model for estimating effect size in the presence of publication bias. *Psychometrika*, 60, 419–435. <http://dx.doi.org/10.1007/BF02294384>
- Vevea, J. L., & Woods, C. M. (2005). Publication bias in research synthesis: Sensitivity analysis using a priori weight functions. *Psychological Methods*, 10, 428–443. <http://dx.doi.org/10.1037/1082-989X.10.4.428>
- Visser, P. S., & Krosnick, J. A. (1998). Development of attitude strength over the life cycle: Surge and decline. *Journal of Personality and Social Psychology*, 75, 1389–1410. <http://dx.doi.org/10.1037/0022-3514.75.6.1389>
- Visser, P. S., Krosnick, J. A., & Lavrakas, P. J. (2000). Survey research. In *Handbook of research methods in social and personality psychology* (pp. 223–252). New York, NY: Cambridge University Press.
- von Helversen, B., Gendolla, G. H. E., Winkielman, P., & Schmidt, R. E. (2008). Exploring the hardship of ease: Subjective and objective effort in the ease-of-processing paradigm. *Motivation and Emotion*, 32, 1–10. <http://dx.doi.org/10.1007/s11031-008-9080-6>
- Weick, M., & Guinote, A. (2008). When subjective experiences matter: Power increases reliance on the ease of retrieval. *Journal of Personality and Social Psychology*, 94, 956–970. <http://dx.doi.org/10.1037/0022-3514.94.6.956>
- Weingarten, E., & Hutchinson, J. W. (2018). Does ease mediate the ease-of-retrieval effect? A meta-analysis. *Psychological Bulletin*, 144, 227–283. <http://dx.doi.org/10.1037/bul0000122>
- Wells, W. D. (1993). Discovery-oriented consumer research. *Journal of Consumer Research*, 19, 489–504. <http://dx.doi.org/10.1086/209318>
- Wikipedia Contributors. (2018). *MSN TV*. Retrieved from [https://en.wikipedia.org/w/index.php?title=MSN\\_TV&oldid=856443875](https://en.wikipedia.org/w/index.php?title=MSN_TV&oldid=856443875)

Wilson, T. D., Aronson, E., & Carlsmith, K. M. (2010). Experimentation in social psychology. In D. T. Gilbert, S. T. Fiske, & G. Lindzey (Eds.), *Handbook of social psychology* (5th ed., Vol. 1, pp. 51–81). Hoboken, NJ: Wiley.

Yeager, D. S., Krosnick, J. A., Chang, L., Javitz, H. S., Levendusky, M. S., Simpson, A., & Wang, R. (2011). Comparing the accuracy of RDD telephone surveys and internet surveys conducted with probability and

non-probability samples. *Public Opinion Quarterly*, 75, 709–747. <http://dx.doi.org/10.1093/poq/nfr020>

Received October 30, 2018

Revision received June 6, 2019

Accepted June 16, 2019 ■