General Article

Why Do Some Psychology Researchers Resist Adopting Proposed Reforms to Research Practices? A Description of Researchers' Rationales

Anthony N. Washburn[1], Brittany E. Hanson[1], Matt Motyl[1],
Linda J. Skitka [1], Caitlyn Yantis[1], Kendal M. Wong[1],
Jiaqing Sun[2], J. P. Prims[1], Allison B. Mueller[1], Zachary J. Melton[1],
and Timothy S. Carsel[1]

[1]Department of Psychology and [2]Department of Managerial Studies, University of Illinois at Chicago

Abstract

In response to the replication crisis, many psychologists recommended that the field adopt several proposed reforms to research practices, such as preregistration, to make research more replicable. However, how researchers have received these proposals is not well known because, to our knowledge, no systematic investigation into use of these reforms has been conducted. We wanted to learn about the rationales researchers have for not adopting the proposed reforms. We analyzed survey data of 1,035 researchers in social and personality psychology who were asked to indicate whether they thought it was acceptable to not follow four specific proposed reforms and to explain their reasoning when they thought it was acceptable to not adopt these reforms. The four reforms were preregistering hypotheses and methods, making data publicly available online, conducting formal power analyses, and reporting effect sizes. Our results suggest that (a) researchers have adopted some of the proposed reforms (e.g., reporting effect sizes) more than others (e.g., preregistering studies) and (b) rationales for not adopting them reflect a need for more discussion and education about their utility and feasibility.

In a recent string of heavily disseminated findings, several published results in cognitive, social, and personality psychology failed to be replicated (Open Science Collaboration, 2015). One explanation for this replication crisis is that many researchers engage in *questionable research practices* (Simmons, Nelson, & Simonsohn, 2011). In recent surveys, many psychologists admitted to deliberately engaging in questionable research practices, including stopping data collection upon finding significant results and reporting unexpected findings as hypothesized a priori (i.e., hypothesizing after the results are known, or HARKing; Fiedler & Schwarz, 2016; John, Loewenstein, & Prelec, 2012; Motyl et al., 2017). As a response to these revelations and others regarding the tenuous nature of replicability in psychological science,

leaders in the field created the Society for Personality and Social Psychology (SPSP) Presidential Task Force on Publication and Research Practices to promote, among other goals, researchers' use of proposed reforms to research practices (PRRPs), such as conducting formal power analyses and reporting effect sizes (Funder et al., 2014). Similarly, the Center for Open Science convened a committee of scholars to establish industry standards that would help promote transparency and openness in the social sciences, including preregistering study

Corresponding Author:
Anthony N. Washburn, University of Illinois at Chicago, Psychology, 1007 W. Harrison St., Chicago, IL 60607-7101
E-mail: awashbu1@uic.edu

hypotheses and making data publicly available online (the Transparency and Openness Promotion, or TOP, Guidelines; Nosek et al., 2015).

What remains unclear, however, is the extent to which researchers are implementing these standards. Initial evidence suggests that although some have enthusiastically embraced PRRPs (e.g., Center for Open Science, 2013; Eich, 2014; Schweinsberg et al., 2016), others have argued for caution with regard to making certain practices, such as preregistration and making data publicly available online, prerequisites for publication (e.g., Finkel, Eastwick, & Reis, 2015). For example, some researchers have argued that preregistering hypotheses and methods can stifle scientific discovery and that the proposed reforms may be less applicable to replications of large field studies or to studies involving difficult-to-recruit populations than to short-term experimental studies (e.g., Goldin-Meadow, 2016a, 2016b; Scott, 2013). Researchers may therefore resist adopting PRRPs for myriad reasons, ranging anywhere from concerns about practical implementation (e.g., institutional-review-board, or IRB, restrictions on sharing potentially identifiable information) to concerns about the utility and need for such practices.

Our goal in conducting the research reported here was to categorize and synthesize researchers' explanations for abstaining from using these practices. In a previous report on a survey of researchers (Motyl et al., 2017) we provided some preliminary examples of justifications for not using PRRPs, but that work was not meant to provide a complete summary of the most frequent reasons why researchers do not always use these practices. Without further analysis, we could not be sure whether the examples we provided in that report are representative of where the field stands as a whole or whether there are other rationales that might be critical in understanding how to proceed in discussions of reforms to research practices. In the current research, we used inductive content coding to arrange all of the participant-generated justifications into broader categories. We then were able to take a closer look at how different categories of justifications are more or less prevalent among actual researchers in the field. This research constitutes an important additional step in the ongoing conversation about how to remediate the replication crisis.

Disclosures

Data, materials, and online resources

The survey materials are available on the Open Science Framework (https://osf.io/c4xtb). Because of IRB restrictions, we are not able to publicly share any data

from this project. A preprint version of Motyl et al. (2017) is available at https://psyarxiv.com/qxaby/.

Measures

We report how we determined our sample size, all data exclusions, and all measures in the study.

Ethical approval

The research reported in this article was conducted under the approval of the institutional review board of the Office for the Protection of Research Subjects at the University of Illinois at Chicago (Protocol 2015-0946).

Method

The data for this study come from a larger survey about research practices and perceptions of the status of social and personality science (see Motyl et al., 2017, for a description of the full data set). The self-reported prevalences of use of the PRRPs were originally published in Motyl et al. (2017). However, results for responses to questions on the acceptability of research practices and results of the inductive content coding of respondents' justifications form the substance of the current research and have not been previously published. We present the self-reported prevalence of each PRRP alongside our new findings to provide a clear picture of how respondents were thinking about these research practices.

Sampling and participants

We e-mailed the members of the SPSP (6,172 members), the European Society for Social Psychology (1,200 members), and the Society of Australasian Social Psychologists (166 members) to solicit their participation in our study.[1] We attempted to reach as many social and personality psychologists (faculty, postdocs, and graduate students) as possible, and our sample consisted of those who replied to our solicitation (see Motyl et al., 2017, for more detailed information about response rates). All survey responses were collected between September 29, 2015, and January 11, 2016. In total 1,414 people opened the survey, and the number of responses to the items relevant to the current study ranged from 1,016 to 1,035.

Most of the respondents who provided background information identified themselves primarily as social psychologists (78%) or personality psychologists (8%). Our sample was 49% male and 47% female (the remainder preferred not to indicate their gender; SPSP's overall gender breakdown is 40% male and 55% female—see

SPSP, 2016). Thirty-two percent of our sample were graduate students, 14% were non-tenure-track degree holders (e.g., adjuncts, postdocs), 19% were assistant professors, 15% were associate professors, and 20% had the rank of full professor or above. Sixty-nine percent of our sample were affiliated with a public university, 30% were affiliated with a private university, and 1% did not have a university affiliation.

Measures

Prevalence of PRRPs. Among other questions, our survey asked researchers how often they used four PRRPs: preregistering studies, making data publicly available online, conducting formal a priori power analyses, and reporting effect sizes. Responses were made on a 5-point scale (*never, rarely, sometimes, often, always*).

Acceptability of not using PRRPs. We also asked the researchers to indicate the extent to which they believed that not engaging in each of these PRRPs was acceptable. The 7-point response scales for these items ranged from *very unacceptable* to *uncertain* (the midpoint) to *very acceptable*.

Rationale for not using PRRPs. Researchers who responded that not using a given PRRP was "slightly," "moderately," or "very" acceptable or who indicated that they were uncertain were asked to write a rationale for their answer in a provided text box (e.g., "When is NOT reporting effect sizes an acceptable research practice?").[2]

Content coding

We employed an inductive coding strategy to identify themes in the rationales for not using each of the PRRPs (e.g., Boyatzis, 1998; Epstein & Martin, 2005). First, as a group, we identified and discussed low-level themes based on similarities across rationales. Similar themes were collapsed together into higher-order categories that were specific enough to be distinct yet broad enough to capture sets of ideas contained in the lower-level themes (e.g., Joffe & Yardley, 2004). The number of discrete themes (e.g., effect sizes do not affect research validity, making data publicly available could lead to misuse of the data) ranged from 7 to 15 across the four PRRPs (see Tables 1 through 4 for lists of the substantive themes used to code the rationales). Once the themes were established, two different authors coded the theme for each rationale for each practice. Because the researchers in our sample sometimes gave multiple rationales for not engaging in a practice, only the first rationale given was coded (note that some responses were uncodable, and others were coded as snark, rather than as belonging to a meaningful theme). Interrater reliability in theme coding was acceptable for each practice (preregistration: κ = .65; making data publicly available: κ = .66; conducting formal power analyses: κ = .63; reporting effect sizes: κ = .66; see Landis & Koch, 1977, for Cohen's κ guidelines). Any disagreements in coding were discussed and resolved by the first and second authors.

Results

Prevalence of the PRRPs

As originally reported (Motyl et al., 2017), self-reported use of each of the PRRPs varied considerably. Almost all (99%) the researchers indicated that they had reported effect sizes, most (87%) said that they had conducted formal power analyses, just over half (56%) reported making their data publicly available online, and few (27%) indicated that they had preregistered studies. Not surprisingly, acceptance of not using these PRRPs was inversely related to the prevalence of these practices (e.g., respondents were most accepting of not preregistering studies). Although these data do not speak directly to day-to-day usage of these practices, they do at least provide a glimpse of how much researchers are taking these practices into consideration when designing or conducting their research. We next examined the reported acceptability of and rationales for not using each PRRP.

Preregistering studies

A majority of the researchers (85.38% of those who answered this item; *n* = 876) indicated that not preregistering studies was acceptable or reported being uncertain about the acceptability of this practice. Of these 876 researchers, 73.97% provided a rationale for their answer.

The most common rationale was that not preregistering studies is acceptable when conducting exploratory or descriptive research (e.g., conducting pilot studies, analyzing large-scale data sets, conducting qualitative research, and conducting secondary data analyses). The second most common rationale for not preregistering studies was that there is currently no requirement for preregistration or the practice is not necessary. For example, many researchers stated that preregistration is not the norm and does nothing to increase the validity of the research. The third most common rationale for not preregistering studies was that preregistration is unnecessary for honest, ethical researchers—in other words, that if researchers are honest in their methods and results, preregistering adds nothing to the integrity

Table 1. Frequencies and Examples of Rationales for Not Preregistering Studies

Rationale	Response rate (%)
Preregistration is not needed for exploratory research (e.g., it is not needed for pilot studies, descriptive research, or secondary data analysis)	33.33
There is no current requirement for preregistration, it does not increase validity (e.g., there is no incentive for preregistration, it is not common practice)	20.52
Preregistration is not needed if researchers are honest (e.g., preregistration creates an atmosphere of distrust, most researchers are trustworthy)	7.87
There are no resources for preregistration, researchers do not know how to preregister studies (e.g., guidelines are unclear, journals do not provide a way to preregister studies)	7.87
Preregistration is costly, is an unnecessary burden (e.g., it takes too much time, there is too much bureaucracy)	6.48
Preregistration constrains researcher degrees of freedom, ignores serendipitous findings (e.g., many interesting findings are unexpected, let the data speak)	5.40
People can still engage in unethical research practices (e.g., researchers can preregister studies after they are run and can still engage in questionable research practices)	3.24
Preregistration may result in being scooped (e.g., other researchers may run preregistered studies)	2.31
Hypotheses are "registered" elsewhere (e.g., hypotheses are registered by being included in IRB protocols or grant applications)	2.01
Other (e.g., preregistering will bias the research)	1.70
Preregistration is not needed when hypotheses are obvious or when conducting programmatic research (e.g., when theory clearly points to specific hypotheses)	1.23
Not preregistering is fine as long as researchers replicate their findings (e.g., it is not needed when results are consistent or robust)	1.08
Preregistration is not needed when studies are conducted as part of a course or for research training (e.g., graduate students' projects are intended just to help students learn the process)	1.08
Snark (e.g., preregistration is silly or idiotic)	0.93
Preregistration creates problems when reviewers and editors evaluate manuscripts (e.g., unexpected findings will be stigmatized)	0.93
Requiring preregistration overvalues confirmatory research (e.g., unregistered findings are not less valid than registered findings)	0.62

Note: A response was considered uncodable when it indicated that the researcher was unsure about the acceptability of not preregistering studies or misunderstood the question; 3.40% of the rationales for not preregistering studies were uncodable. IRB = institutional review board.

of the research. The next few most common rationales reflected uncertainty about how to preregister studies in the first place and perceptions that preregistration is burdensome and restrictive (see Table 1 for frequencies and example rationales for all the themes).

Making data publicly available online

Seven hundred twenty-six researchers (70.35% of those who answered this item) indicated that not making data available online was acceptable or reported being uncertain about the acceptability of this practice. Of these 726 researchers, 74.38% provided a rationale for their response.

The most common rationale for not making data publicly available online referred to participants' confidentiality. Researchers were particularly concerned about protecting sensitive data (e.g., proprietary organizational information) and easily identifiable data (e.g., data from couples in close-relationships research). The second most common rationale revolved around intellectual-property issues. For example, the researchers in our

sample reported that they were worried about making data available to other researchers to use before they themselves had analyzed all the publishable data (e.g., concern about getting scooped). The third most common rationale for not making data publicly available reflected existing data-sharing norms. Many respondents stated that they would share their data with other researchers upon request, which justified not sharing the data publicly. Most of the remaining rationales reflected perceptions that making data publicly available is not necessary, that there is no standard way to do so, and that there are too many institutional and other restrictions (e.g., IRB issues, participants' consent; see Table 2 for frequencies and example rationales for all the themes).

Conducting formal a priori power analyses

Six hundred fourteen researchers (59.90% of those who answered this item) indicated that not conducting an a priori power analysis was acceptable or reported being

Table 2. Frequencies and Examples of Rationales for Not Making Data Publicly Available Online

Rationale	Response rate (%)
Making data publicly available raises issues concerning participants' confidentiality (e.g., participants' identity needs to be protected when data are sensitive)	27.22
Making data publicly available invites intellectual-property issues (e.g., the researcher may not be finished analyzing all the data that could be published in future manuscripts)	16.48
Researchers will share data upon request (e.g., data do not need to be publicly available to be open to sharing)	15.37
This proposed research-practice reform is not needed (e.g., not sharing data is always okay)	8.52
Restrictions prevent data sharing (e.g., IRBs will not allow data sharing, participants did not consent)	8.33
The logistics of how to make data publicly available are not well known (e.g., making data publicly available is not the norm, it is unclear where to post data online)	5.93
Making data public is too costly (e.g., it is too time-consuming to prepare data for an online repository)	4.07
Other (e.g., data may be uninformative, some researchers admit to being lazy)	3.70
Making data publicly available could lead to misuse of the data (e.g., data may be too complex for non-content experts to analyze)	2.96
This proposed research-practice reform does not actually improve science (e.g., data quality cannot be inferred from data posted online)	2.78
Snark (e.g., data may be kept private to hide misdeeds)	0.56

Note: A response was considered uncodable when it indicated that the researcher was unsure about the acceptability of not making data publicly available or misunderstood the question; 4.07% of the rationales for not making data publicly available were uncodable. IRB = institutional review board.

uncertain about the acceptability of this practice. Of these 614 researchers, 67.92% provided a rationale for their answer.

The most common rationale for not conducting power analyses was that there is not always a basis for estimating the true effect size in a study. For example, some researchers pointed out that power analyses are not possible when the effect sizes are unknowable (e.g., in new research paradigms or exploratory settings). The second most common theme was that experience makes it possible to guess the appropriate sample size for a study without conducting a formal power analysis. For example, many respondents stated that they already know how large their sample needs to be to find effects in programmatic research. The third most common category of responses consisted of rationales that were uncodable. The majority of these rationales reflected uncertainty about why it is acceptable to not conduct power analyses or misunderstanding of the question. The next most common rationale for not conducting

Table 3. Frequencies and Examples of Rationales for Not Conducting Formal Power Analyses

Rationale	Response rate (%)
There may be no basis for estimating effect size, power analyses are not needed for exploratory research (e.g., effect sizes are not known in the case of new research paradigms)	30.94
Power analyses are not needed when researchers have experience or can guess the appropriate sample size without formal analyses (e.g., in programmatic research, researchers know how big the sample should be)	15.11
Limited resources or other constraints on increasing sample sizes make power analyses unnecessary (e.g., sample size may be limited by the size of the participant pool)	10.55
Power analyses are not normative or required (e.g., journals do not require power analyses)	8.39
Other (e.g., researchers may intentionally run underpowered studies)	8.39
Power analyses are not needed when planned or actual sample sizes are huge (e.g., some sample sizes are large enough that no analysis is needed to determine that even small effects can be detected)	5.52
Power analyses are ridiculous or not valid (e.g., researchers can arbitrarily pick the effect size for a power analysis)	5.04
Power analyses are not useful for secondary data analyses and other cases in which researchers have no control over sample size (e.g., sample size may be unchangeable)	1.20
Snark (e.g., power analyses are not needed when researchers do not care about quality science)	0.72
Qualitative research does not require power analyses (e.g., formal effect sizes are hard to determine for qualitative research)	0.48

Note: A response was considered uncodable when it indicated that the researcher was unsure about the acceptability of not conducting power analyses or misunderstood the question; 13.67% of the rationales for not conducting power analyses were uncodable.

Table 4. Frequencies and Examples of Rationales for Not Reporting Effect Sizes

Rationale	Response rate (%)
Effect sizes are redundant with other reported statistics or should be calculated by readers (e.g., all the information is there for readers to calculate effect sizes)	25.12
Effect sizes are not important, are uninformative, or are irrelevant to the field or specific research questions (e.g., psychologists are interested in process, not intervention)	21.67
Effect-size reporting is not a standard, norm, or journal requirement (e.g., some journals, reviewers, and editors do not ask for effect sizes)	15.76
Proper methods for calculating effect size are not available or are unclear (e.g., it is difficult to calculate effect sizes for multilevel analyses)	11.82
Effect sizes do not affect research validity (e.g., effect sizes say nothing about the integrity of the research)	9.85
Other (e.g., researchers may not be sure how to interpret effect sizes)	3.94

Note: A response was considered uncodable when it indicated that the researcher was unsure about the acceptability of not reporting effect sizes or misunderstood the question; 11.82% of the rationales for not reporting effect sizes were uncodable.

power analyses reflected concerns about limited participant pools and resources. Many researchers stated that their participant pool was finite and a power analysis might indicate the need for more participants than it would be possible to recruit. The remaining rationales included the idea that high power is always achieved by having large enough samples and the fact that researchers do not always have control over sample sizes (e.g., when they are conducting secondary data analyses on publicly available data sets; see Table 3 for frequencies and example rationales for all the themes).

Reporting effect sizes

Three hundred sixty-two researchers (35.63% of those who answered this item) indicated that not reporting effect sizes was acceptable or reported being uncertain about the acceptability of this practice. Of these 362 researchers, 56.08% provided a rationale for their answer.

The most common rationale for not reporting effect sizes was the belief that they are redundant with other reported statistics and can be inferred by readers. For example, many respondents stated that effect sizes can be computed directly from the test statistics (e.g., *F* or *t* tests and degrees of freedom) and therefore do not need to be reported explicitly. The second most common rationale was that effect sizes are unimportant, uninformative, or irrelevant for certain types of research questions. Some of the researchers reported that psychologists are more interested in process than in intervention and that effect sizes are therefore unimportant and potentially misleading. Others noted that manipulations and measures, particularly in social psychology, are arbitrary and that effect sizes are therefore similarly arbitrary and uninformative because they are not measuring the size of a "real life" effect. The third most common rationale was that reporting effect sizes is not

standard or the norm and is not often required by journals. Finally, some respondents argued that methods for calculating effect sizes are not readily available or are unclear and that effect sizes themselves do not speak to the validity of the research (see Table 4 for frequencies and example rationales for all the themes).

Discussion

Our results suggest that the rationales for not adopting PRRPs in psychology are varied. Many researchers' rationales point to possible limitations of the current discussion on research-practice reform. For example, most of the researchers in our sample revealed concerns about how to implement PRRPs under specific conditions (e.g., preregistering exploratory research, adhering to IRB restrictions on data sharing). Concerns about how to implement PRRPs indicate that the discussion needs to more specifically address how to apply these practices in the case of research that is not traditional laboratory research, such as field studies and research on sensitive populations (see Finkel et al., 2015, for a discussion). Furthermore, guidelines for how to incorporate these practices into IRB protocols and consent documents may help researchers participate in open science while still respecting participants' privacy, particularly as federal regulations change (e.g., Jaschik, 2017). In addition to these practical considerations, some researchers provided rationales that reflect open questions (e.g., do researchers maintain authorship rights over their publicly available data?) that will require continued discussion to establish field norms.

Some researchers in our sample seemed to be resisting using PRRPs because they thought these practices are not needed or that the costs associated with them create an unnecessary burden on researchers (e.g., effect sizes are redundant, preregistering is another

bureaucratic hurdle). Thus, efforts should focus on combating resistance through clarifying the utility and benefit of using such practices. For example, even if researchers are honest in their endeavors, preregistration (and other PRRPs) may still protect them from unwittingly engaging in questionable research practices (e.g., HARKing; Kerr, 1998). In addition, our results indicate that increasing the ease of use of preregistration platforms (and therefore reducing effort and costs) may encourage more widespread adoption of this practice.

Finally, some rationales indicated that there is still a great deal of confusion about what PRRPs are and how one goes about implementing them. Some of the researchers in our sample indicated that they were unaware that preregistering studies and making data available online, for example, were considered best practices or did not know how or where to preregister studies or make their data publicly available. Therefore, it would be beneficial to more broadly inform researchers in psychology about open-science initiatives that provide platforms for storing preregistrations, study materials, and data online (e.g., Open Science Framework, the Dataverse Project, AsPredicted.org). These types of responses suggest that the field still has a way to go to educate all researchers—including those who may not be able to routinely attend conferences, for example—about still quite-new platforms for preregistering studies and sharing data.

In summary, improving psychological science with PRRPs will require not only breaking down perceived barriers for researchers who want to implement these reforms but do not think they can or, simply, are not sure how to do so, but also addressing the fact that some researchers do not see a need to reform research practices in the first place.

Action Editor

Simine Vazire served as action editor for this article.

Author Contributions

A. N. Washburn and B. E. Hanson drafted the manuscript. B. E. Hanson and A. N. Washburn analyzed the data. M. Motyl conceived of the idea for the original survey. L. J. Skitka and M. Motyl developed the survey. All the authors developed the coding schemes and coded the justifications. M. Motyl and L. J. Skitka read, revised, and provided critical feedback on the manuscript. All the authors approved the final version of the manuscript for submission. Starting with C. Yantis, the authors are listed in reverse alphabetical order.

ORCID iD

Linda J. Skitka (iD) https://orcid.org/0000-0002-7078-0650

Declaration of Conflicting Interests

The author(s) declared that there were no conflicts of interest with respect to the authorship or the publication of this article.

Open Practices

All materials have been made publicly available via the Open Science Framework and can be accessed at https://osf.io/c4xtb/. The complete Open Practices Disclosure for this article can be found at http://journals.sagepub.com/doi/suppl/10.1177/2515245918757427. This article has received the badge for Open Materials. More information about the Open Practices badges can be found at http://www.psychological science.org/publications/badges.

Notes

1. We also contacted the Asian Association of Social Psychology, but they declined to distribute the survey to their members.
2. Because of a programming error, 71 researchers who said that not reporting effect sizes was "moderately acceptable" were not given a text box to write a rationale.

References

Boyatzis, R. E. (1998). *Transforming qualitative information.* London, England: Sage.

Center for Open Science. (2013). *Badges to acknowledge open practices.* Retrieved from https://openscienceframework.org/project/TVyXZ/

Eich, E. (2014). Business not as usual. *Psychological Science,* *25,* 3–6.

Epstein, L., & Martin, A. D. (2005). Coding variables. In K. Kempf-Leonard (Ed.), *Encyclopedia of social measurement* (Vol. 1, pp. 321–327). Oxford, England: Elsevier Academic Press.

Fiedler, K., & Schwarz, N. (2016). Questionable research practices revisited. *Social Psychological & Personality Science,* *7,* 45–52.

Finkel, E. J., Eastwick, P. W., & Reis, H. T. (2015). Best research practices in psychology: Illustrating epistemological and pragmatic considerations with the case of relationship science. *Journal of Personality and Social Psychology,* *108,* 275–297.

Funder, D. C., Levine, J. M., Mackie, D. M., Morf, C. C., Sansone, C., Vazire, S., & West, S. G. (2014). Improving the dependability of research in personality and social psychology: Recommendations for research and educational practice. *Personality and Social Psychology Review,* *18,* 3–12.

Goldin-Meadow, S. (2016a). Preregistration, replication, and nonexperimental studies. *Observer,* *29*(8), 5, 14.

Goldin-Meadow, S. (2016b). Why preregistration makes me nervous. *Observer,* *29*(7), 1–2.

Jaschik, S. (2017, January 19). New 'common rule' for research. *Inside Higher Ed.* Retrieved from https://www

.insidehighered.com/news/2017/01/19/us-issues-final-version-common-rule-research-involving-humans

Joffe, H., & Yardley, L. (2004). Content and thematic analysis. In D. Marks, & L. Yardley (Eds.), *Research methods in clinical and health psychology* (pp. 56–68). London, England: Sage.

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.

Kerr, N. L. (1998). HARKing: Hypothesizing After the Results are Known. *Personality and Social Psychology Review, 2*, 196–217.

Landis, J. R., & Koch, G. G. (1977). The measurement of observer agreement for categorical data. *Biometrics, 33*, 159–174.

Motyl, M., Demos, A. P., Carsel, T. S., Hanson, B. E., Melton, Z. J., Mueller, A. B., . . . Skitka, L. J. (2017). The state of social and personality science: Rotten to the core, not so bad, getting better, or getting worse? *Journal of Personality and Social Psychology, 113*, 34–58.

Nosek, B. A., Alter, G., Banks, G. C., Borsboom, D., Bowman, S. D., Breckler, S. J., . . . Yarkoni, T. (2015). Promoting an open research culture. *Science, 348*, 1422–1425.

Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science, 349*, 943–951.

Schweinsberg, M., Madan, N., Vianello, M., Sommer, S. A., Jordan, J., Tierney, W., . . . Uhlmann, E. L. (2016). The pipeline project: Pre-publication independent replications of a single laboratory's research pipeline. *Journal of Experimental Social Psychology, 66*, 55–67.

Scott, S. (2013). Pre-registration would put science in chains. *Times Higher Education*. Retrieved from http://www.timeshighereducation.co.uk/comment/opinion/pre-registration-would-put-science-in-chains/2005954.article

Simmons, J. P., Nelson, L. N., & Simonsohn, U. (2011). False-positive psychology: Undisclosed flexibility in data collection and analysis allows presenting anything as significant. *Psychological Science, 22*, 1359–1366.

Society for Personality and Social Psychology. (2016). *Membership diversity statistics, SPSP, Inc. (12/31/2016)*. Retrieved from http://spsp.org/sites/default/files/Membership%20Diversity%20December%202016.pdf