Running head: Gender bias by S.T.E.M. faculty?

This is a stage 1 registered replication report that has been given in-principle acceptance at Meta-Psychology prior to data collection. The editorial process, including decision letters and reviews, can be accessed at:

<http://tinyurl.com/mp-submissions>

Anyone can contribute with Open Peer Review through hypothes.is directly on this preprint. However, note that the IPA decision means that your peer review cannot change the data collection as it's already under way.

Are STEM Faculty Biased Against Female Applicants?

A Robust Replication and Extension of Moss-Racusin and Colleagues (2012)

Nathan Honeycutt, Lee Jussim, and Akeela Careem

Rutgers University

Neil A. Lewis, Jr.

Cornell University

NOTE: this is a registered replication report shared as a pre-print for the purposes of pre-registration and open review. Originally shared 8/18/20, revised 4/29/21, revised again 5/24/21 to reflect having obtained in-principle acceptance (IPA).

Acknowledgements: We would like to thank Dr. Corinne Moss-Racusin for her extensive feedback on our methods and materials, which have helped maximize the similarity of our replication to the original study.

Abstract

In a seminal study investigating gender bias in academic science, Moss-Racusin et al. (2012) found bias against female lab manager applicants with respect to competence, hireability, mentoring, and salary conferral. This topic will be revisited through three studies--one direct replication, one extension, and a meta-analysis. The present set of studies, all using the same methods and materials as the original, will sample from a larger and broader pool of faculty, collectively representing the breadth of STEM disciplines at R1 universities across the United States. The proposed studies will have high power to detect smaller effect sizes than in the original. Furthermore, a priori criteria have been explicitly articulated for what will be accepted as hypothesis confirmation and replication, and for evidence of gender bias. The pre-registered data analysis plans will provide a strong test assessing the replicability of Moss-Racusin et al. (2012)'s finding that STEM faculty were biased against women who applied for a lab manager position.

Keywords: Gender bias, STEM, diversity, women in science

Are STEM Faculty Biased Against Female Applicants?

A Robust Replication and Extension of Moss-Racusin and Colleagues (2012)

How prevalent is gender bias against women in academic science, technology, engineering, and mathematics (i.e., “STEM”) disciplines? This critical topic has been at the forefront of many efforts and initiatives aimed to understand why women are underrepresented in STEM disciplines, and to find ways to mitigate it. At the most basic level, gaps (i.e., differences in the proportion of men versus women) have been a critical piece of evidence when discussing gender bias. Most population data consistently indicate that women are underrepresented and men overrepresented in STEM, particularly within engineering, computer, and physical science STEM fields (e.g., Funk & Parker, 2018; Handelsman et al., 2005). With gaps as a starting point, experimental investigations on bias in hiring, promotion, awards, publications, studies on the experiences of women in STEM disciplines, and studies on motivations (or hinderances) to pursuing a STEM career have continued to provide additional and updated information on the representation and experience of women in STEM (see the following section for many such important studies and commentaries).

The present research takes a specific approach to understanding these issues by providing an updated experimental investigation of gender bias in academic STEM. Through the framework of a robust replication and extension of an influential experimental investigation into bias against women in academic science disciplines, we will examine if faculty evaluate men more favorably than women along the dimensions of competence, hireability, salary conferral, and willingness to mentor.

Evidence Suggesting Gender Bias in Academic STEM

A vast literature addresses gender bias in academia (Beaudry & Larivière, 2016; Brown & Goh, 2016; Budden et al., 2008; Carter, Smith, & Osteen, 2017; Ceci et al., 2014; Ceci &

Williams, 2010; Ceci & Williams, 2011; Dion, Sumner, & Mitchell, 2018; Eagly & Miller, 2016; Eaton, Saunders, Jacobson, & West, 2020; Geraci, Balsis, & Busch, 2015; Kaatz, Gutierrez, & Carnes, 2014; Knobloch-Westerwick & Glynn, 2013; Knobloch-Westerwick, Glynn, & Hoge, 2013; Lincoln, Pincus, Koster, & Leboy, 2012; Madera, Hebl, & Martin, 2009; Monroe, Ozyurt, Wrigley, & Alexander, 2008; Milkman, Akinola, & Chugh, 2012; Moss-Racusin, Dovidio, Brescoll, Graham, & Handelsman, 2012; Nitttrouer et al., 2018; Reuben, Sapienza, & Zingales, 2014; Roberts & Verhoef, 2016; Robnett, 2016; Rossiter, 1993; Sheltzer & Smith, 2014; Steele, James, & Barnett, 2002; Steinpreis, Anders, & Ritzke, 1999; Teele & Thelen, 2017; Trix & Psenka, 2003; Valian, 2004; van der Lee & Ellemers, 2015; Wai et al., 2010; Wenneras & Wold, 1997; West, Jacquet, King, Correll, & Bergstrom, 2013; Williams, Phillips, & Hall, 2016).

Although it is beyond the scope of the present research to provide a comprehensive review of that literature, we briefly discuss some of the findings most tightly related to the present investigation.

Several studies have revealed evidence of gender bias in the academic hiring process. In an audit-type study, Moss-Racusin et al. (2012) found that science faculty members displayed gender bias against women when asked to rate an applicant for a lab manager position. Participants all read the same application material, with just the name/gender manipulated. Both male and female faculty members rated the female applicant as less competent and less hireable than the male applicant, and offered less career mentoring and a lower starting salary to the female applicant (Moss-Racusin et al., 2012).

A similar audit-type study by Eaton et al. (2020) investigated gender and racial bias among physics and biology faculty towards post-doctoral candidates applying for a tenure-track assistant professor position. Faculty participants read the applicant's CV, where gender and race were manipulated using the applicant's name. Physics faculty displayed a gender bias such that

male applicants were rated higher on competence and hireability than the female applicants.

Across departments, women received higher likeability ratings than men.

Another audit-type study revealed that academic psychologists were more likely to hire a male job applicant than an identical female applicant, and they gave the male applicant higher evaluations on research, teaching, and service (Steinpreis et al., 1999); however, participants did not display gender bias when voting to tenure identical male and female applicants (Steinpreis et al., 1999). Other research has found that, among biology labs at many U.S. universities, male faculty members had fewer female graduate students and postdocs than female faculty members (Sheltzer & Smith, 2014). In another study of hiring decisions, student participants were asked to “hire” one person from a pair of candidates to perform an arithmetic task. The task was one that both men and women could perform equally well. Both male and female participants displayed bias such that they were more likely to hire a man, even though they had no information about the candidates other than their appearance, and therefore their gender (Reuben, Sapienza, & Zingales, 2014). When participants were informed of how well the candidates did on the task previously, the gender bias was reduced (although it was still present).

Other research has revealed that women (and minorities) appear to be subject to more bias when people are making decisions about distant future events, rather than near-future events (Milkman, Akinola, & Chugh, 2012). In this study, faculty members were sent emails from a prospective doctoral student asking for a meeting. The student’s name was manipulated to indicate race and gender, and the student asked for a meeting either the same day or in one week. When the meeting was in one week, Caucasian males received more acceptance emails and quicker responses. However, there was no bias when the meeting was requested for the same day (Milkman et al., 2012).

Researchers have also studied self-reported experiences from women to shed more light on gender bias in academia. In a study examining self-reported experiences of gender bias in STEM from women at the high school, undergraduate, and graduate level, 61% of participants reported experiencing bias in the past year (Robnett, 2016), and undergraduate women in math-intensive fields reported experiencing more gender bias than others. Similarly, women reported experiencing more bias in male-dominated (than in female-dominated) areas of math, science and engineering (Steele, James, & Barnett, 2002). In addition, when using interviews and surveys to investigate female scientists' experiences of bias, researchers have found that a majority of female scientists' surveyed reported experiencing gender bias (Williams, Phillips, & Hall, 2016). Minority women also self-reported experiences of racial discrimination, such as having to prove one's self repeatedly, and being pressured to take on traditionally feminine roles (Williams et al., 2016).

In investigations of gender bias in forms other than hiring, several studies have found that women also experience bias in the academic peer review process. A Swedish study on the peer-review process for funding applications in the biomedical field (Wenneras & Wold, 1997) found that female applicants received lower ratings than male applicants on scientific competence, quality of proposed methods, and relevance of the proposal. Notably, women were rated lower on scientific competence than men who demonstrated the same productivity level; the authors estimated that women had to be 2.5 times more productive than men to receive the same competence score (Wenneras & Wold, 1997). Similarly, in the Netherlands, an analysis of research funding decisions by a science organization revealed gender bias against women in application evaluations and success rates. Women received lower evaluations than men on the quality of the researcher (but not on quality of proposal), and gender differences were also

identified in the language used by evaluators on instruction and evaluation sheets (van der Lee & Ellemers, 2015).

In other investigations of gender bias, an ecology journal discovered that after they introduced double-blind review, the proportion of female first author papers increased by approximately 7.9% (Budden et al., 2008). In comparison, a similar ecology journal using single-blind review did not see a significant increase in female first-author representation in the same time period (Budden et al., 2008). Similarly, the EvoLang academic conference found that introducing double-blind review reduced gender bias against female authors, as compared to previously used single-blind review (Roberts & Verhoef, 2016).

In another study related to the peer-review process, science communications graduate students were asked to rate conference abstracts with the author names varied to manipulate gender (Knobloch-Westerwick, Glynn, & Huge, 2013). Abstracts received higher scientific quality ratings if participants thought they were written by male authors as opposed to female authors. This bias was especially present if the abstract topic was in a male-typed research area (Knobloch-Westerwick et al., 2013).

Gender gaps have also been found in author citation patterns, with female authors tending to receive fewer citations of their scholarly work than male authors (Beaudry & Larivière, 2016). A similar gap was found in social and personality psychology, a field in which women outnumber men according to membership statistics from the Society for Personality and Social Psychology (Brown & Goh, 2016). But despite this emerging membership trend, the authors found that men receive more awards than women (Brown & Goh, 2016). Their results mirror findings in other areas, such as an analysis of awards data from 13 STEM societies revealing that men received more awards than women, and this number was higher than expected based on the proportion of male to female nominees in the pool (Lincoln et al., 2012).

Evidence Suggesting Lack of Gender Bias in Academic STEM

Alternatively, there is some research that has failed to obtain evidence of gender bias in academia, and even some papers that have obtained evidence of reverse bias (i.e., biases favoring women over men). In a series of five experiments investigating gender bias in tenure-track faculty hiring for STEM fields (both math-intensive and non-math intensive fields), involving 873 tenure-track faculty participants, faculty displayed a 2:1 bias in favor of female applicants over identically qualified male applicants, except male economists who displayed no gender bias (Williams & Ceci, 2015). In other words, roughly 66% of participants ranked the female applicant first.

Some studies have also found no evidence of gender bias in the peer-review process. In a review of some ecology and evolution journals over a one-year period, Tregenza (2002) did not find evidence of gender bias (towards the first author) in the manuscript acceptance process and in citations. A similar analysis of papers submitted to a biology journal also failed to reveal evidence of gender bias in the current single-blind peer review process (Primack, Ellwood, Miller-Rushing, Marrs, & Mulligan, 2009). Further, other researchers have contested findings such as those by Budden et al. (2008), which concluded that female first-author papers had increased significantly after an Ecology journal introduced double-blind review. Upon reanalysis of the data, other researchers found that female authorship increased at about the same rate at journals *not* adopting double-blind review (Webb, O'Hara, & Freckleton, 2008). Therefore, the increase at journals adopting double-blind review was unlikely to have resulted from adoption of double-blind review, suggesting that the single-blind review process was not biased against female authors.

Contrary to prior findings, some research has failed to find evidence of traditional gender bias by reviewers for several behavioral journals. Notably, Lloyd (1990) found that male

reviewers did not display gender bias, but female reviewers displayed gender bias in favor of women when accepting manuscripts. Further, another study found that female scientists display in-group bias, attributing higher levels of scientific traits to female scientists than to male scientists (Veldkamp, Hartgerink, Assen, & Wicherts, 2017).

Other research has also failed to identify gender bias in the peer review process, such as a 2009 meta-analysis which found no gender bias in the peer-review process for grants (Marsh, Bornmann, Mutz, Daniel, & O'Mara, 2009). Additionally, several studies investigating the NIH grant funding process revealed equal success rates for men and women in almost all types of grants (Ley & Hamilton, 2008; Pohlhaus, Jiang, Wagner, Schaffer, & Pinn, 2011). An experimental study in this domain found no evidence of gender bias in the NIH initial evaluation process for R01 grants (Forscher, Cox, Brauer, & Devine, 2018). Further, research by van der Lee & Ellemers (2015), which revealed evidence of gender bias in research funding decisions, was reanalyzed by Albers (2015) who found that the original analysis suffered from Simpson's Paradox (i.e., although women did receive fewer awards, they also applied more often in fields where funding was more difficult to obtain). The reanalysis of the data by Albers (2015) revealed no evidence of gender bias.

Our review of the literature has revealed, however, a complex pattern of highly varied and sometimes even conflicting findings. For both theoretical and applied reasons it is crucial, therefore, to conduct research in this area in order to maximize the certainty with which clear and justified conclusions can be reached, especially in light of such varying findings. There are many routes by which such certainty can be increased. One is to conduct close replications of highly influential and foundational studies; this is the purpose of the proposed research.

Proposed Replication of Moss-Racusin et al. (2012) Study

The present research seeks to provide an updated focused replication of a highly influential paper demonstrating gender bias in academic hiring (Moss-Racusin et al., 2012). The Moss-Racusin et al. (2012) paper stands as one of the most highly cited (2,625 citations as of April 29, 2021) and discussed (Altmetric attention score of 5601 as of April 29, 2021) empirical studies investigating gender bias in academic science. In the study, tenured/tenure-track science faculty members in biology, chemistry, and physics were asked to rate a qualified, but not overly excellent, applicant for a lab manager position. Participants all read the same application material, but the name/gender on the application was randomly assigned: the ostensible applicant was named either “John” or “Jennifer.”

Both male and female faculty members rated the female applicant as less competent and less hireable than the male applicant. Further, they offered less career mentoring and a lower starting salary to the female applicant. Scores on the Modern Sexism Scale (Swim, Aikin, Hall, & Hunter, 1995) predicted ratings of the female applicant’s (but not the male applicant’s) competence, hireability, and willingness to mentor. These findings led the authors to conclude that science faculty members exhibit gender bias that is likely unintentional and unconscious (Moss-Racusin et al., 2012).

Although highly influential, the study has some important limitations. First, the study included a small sample of 127 faculty participants in a two condition between-subjects experimental design, which fell within the conventional norms for research at that time (Simmons, Nelson, & Simonsohn, 2011). The authors justified this sample size using a power analysis targeting a “modest” effect size (Moss-Racusin et al., 2012, Supplementary Materials, p. 1), but did not explain what this meant, leaving readers to guess that it probably was based on Cohen's (1988) standard of $d=.5$. Their description did not identify how much power they had to

detect what precise effect size, nor which analyses their power analysis applied to; therefore, our replication will be guided by our own power considerations, as there was little to no interpretable guidance provided by Moss-Racusin et al. (2012). More important, however, Moss-Racusin et al. (2012) did not justify or explain why they expected a moderate effect size, given that most effects in social psychology are smaller (Richard, Bond, & Stokes-Zoota, 2003). Scientists have recently learned how critical these methodological limitations are for our inferences (Camerer et al., 2018; Fraley & Vazire, 2014; Nelson, Simmons, & Simonsohn, 2018; Open Science Collaboration, 2015), and thus a robust replication is warranted, particularly given the manner in which broad implications have been drawn from their findings (Chang, 2012).

In the original study, data was collected from six universities in different geographic areas, and the authors stated that the participants were generally representative on demographic characteristics. Nonetheless, sampling from only six universities may have limited the generalizability and representativeness of the findings.

Further, whereas in the original paper conclusions were carefully framed as referring to “academic science” or “science fields” (given they only recruited faculty participants from biology, chemistry, and physics departments), subsequent discussions and citations have not been as careful. Many have extrapolated the Moss-Racusin et al. (2012) findings to STEM fields (i.e., science, technology, engineering, and mathematics) more broadly (e.g., Byars-Winston, 2014; Miller & Wai, 2015)—a conclusion that the original data technically cannot support. Whereas biology, chemistry, and physics (the fields examined in the original study) are certainly a component of STEM, they are a very small subset of STEM fields. In order to reach conclusions about “gender bias by STEM faculty” or “gender bias in STEM” rather than the particular three fields Moss-Racusin et al. (2012) studied, evidence that bears on a wider range of STEM fields is clearly needed. In sum, a larger and broader sample of faculty from the breadth

of STEM disciplines at more than six universities is critical to determine if the key results and implications hold.

Conducting a Direct Replication of Moss-Racusin et al. (2012)

The present study will be a direct replication and extension of Moss-Racusin et al. (2012). A direct replication is necessary for the following reasons:

- Gender bias in STEM is potentially a major obstacle to reducing gender disparities. However, the scientific literature on this topic has produced mixed results, with some studies finding evidence of bias, some finding no bias, and some finding biases favoring women. One important way to address these seemingly conflicting findings is to evaluate their replicability.
- Small sample studies have been at the heart of the Replication Crisis. Because small sample sizes produce notoriously variable and therefore uncertain results, a large sample replication is necessary to reduce uncertainty regarding whatever is found (Asendorpf et al., 2013).
- Moss-Racusin et al. (2012) was a high impact study that scientists give high credibility to, so, in addition to finding out whether it replicates with a much larger sample, we need to know:
 - If the study's results can be independently obtained by other scientists.
 - If the study's results can be obtained in a sample that is broader than the six universities included in the original study.
- Scientists have overgeneralized the results of Moss-Racusin et al. (2012) beyond "science" to "STEM" in the sense that the study only addressed S(cience), and not TEM (technology, engineering, and mathematics). However, even though the original Moss-Racusin et al. (2012) study did not address TEM, it is possible that

what we have characterized here as “overgeneralizations” (to TEM from just S) are nonetheless correct. Of course, it is also possible that the overgeneralizations are incorrect. Therefore, we need to know if the results finding bias against women applying for lab manager positions generalizes beyond “science” to “STEM.”

The purpose of the present study is not to comprehensively answer the general question of whether gender bias does or does not exist in academic STEM. That question, unfortunately, is beyond the scope of any single project, including this one. Gender bias may vary in how it operates in different contexts, different environments, et cetera. Because the topic is so important, replication studies are invaluable. The purpose of the present study, therefore, is to conduct a direct replication of the important and influential Moss-Racusin et al. (2012) study.

Adversarial Collaboration

Finally, the current paper is an adversarial collaboration. Whereas most researchers might like to believe that they are unbiased, the reality is that most (and possibly all) individuals are at least potentially subject to biases. This particular project got its start after two of the authors (Second Author, Fourth Author) debated one another regarding the power and importance of implicit bias (Second Author: lower; Fourth Author: higher). The two other members of the team (Third Author, First Author) also had differing viewpoints regarding the likelihood of successfully replicating Moss-Racusin et al. (2012). Consequently, each member could operate as a skeptical check on manifestations of biases among other members of the collaborative team, thereby maximally limiting the extent to which such biases may undermine the quality of the project. We articulated our differences a priori, as shown in the table below (Table 1).

Table 1. Overview of Gender Bias Predictions				
Sample	First Author's Prediction	Second Author's Prediction	Third Author's Prediction	Fourth Author's Prediction
Study 1, 2, 3, and 4	$-.22 < d < .22$	$-.2 < d < .3$	$.2 < d < .5$	$0 < d < .2$
<p><i>Note 1:</i> Positive effect sizes correspond to “Pro-Male Bias”—a gender bias that favors the male student applicant. Negative effect sizes correspond to “Pro Female Bias”—a gender bias that favors the female student applicant. “Pro Male Bias” would be consistent with Moss-Racusin et al. (2012).</p> <p><i>Note 2:</i> Only one prediction is offered for all authors as predictions were the same across all studies for all authors.</p>				

As can be seen, to capture the uncertainty surrounding point estimates, all four generated a range of expected results (a prediction analog to confidence intervals). We were more than a little surprised to discover that, when compelled to generate predictions for effect sizes, most overlapped considerably. Although the first and third authors' predictions had very little overlap, all other pairwise combinations of predictions had a great deal of overlap. On average, we are expecting some gender bias, albeit considerably less than found in the original study, although three of the authors' predictions include zero as within the range of likely outcomes. This shared expectation for a smaller effect size is, in part, a byproduct of uncertainty surrounding the original study—it was a small scale study, and effect sizes are often overestimated in small scale studies (Ioannidis, 2008). Additionally, this expectation is anchored by expectations for reasonable effect sizes in the social cognition area of psychology, given effect sizes researchers in this domain have been obtaining (e.g., Forscher et al., 2019).

Limitations this Study will not Address

There are a number of issues this study will not address. We are conducting a direct replication, not a conceptual replication. Direct replication is necessary to determine if effects are reliable, and whether labs other than the original can obtain the findings (Nature Human Behaviour Editorial Board, 2018; Simons, 2014). There may be other important ways researchers

could expand and improve Moss-Racusin et al. (2012) that will not be a part of this direct replication. For example, it may be important to improve Moss-Racusin et al. (2012) by using names other than “John” and “Jennifer,” or by taking into account ethnicity, or by considering applicants for positions other than lab manager, or by assessing biases in grading, publication, citations, or grant funding rather than hiring. But, all of these changes are beyond the scope of a direct replication. We concur with Simons (2014, p. 76) that:

1. Direct replication by other laboratories is the best (and possibly the only) believable evidence for the reliability of an effect;
2. The idea that only the originating lab can meaningfully replicate an effect limits the scope of our findings to the point of being uninteresting and unfalsifiable.

Research that addresses all of the limitations of Moss-Racusin et al. (2012) is not only beyond the scope of a direct replication, it would mean that there would need to be so many changes that *we would not actually be performing a direct replication*. Given the value and importance of direct replications, we leave improving on the many limitations to Moss-Racusin et al. (2012) to future research.

Present Study

Given the centrality of Moss-Racusin et al. (2012) to discussions of gender bias in academic STEM fields, we will first conduct a pre-registered direct replication (Study 1) utilizing a much larger sample. This will be followed by a pre-registered extension (Study 2) which will utilize a large and broad sample that draws faculty participants from the breadth of STEM disciplines at universities across the United States. Finally, Study 3 will be a pre-registered meta-analysis, combining data from Studies 1 and 2. Following recommendations from Wicherts et al. (2016, p. 2), we sought to be “specific, precise, and exhaustive” in the plans of our studies.

STUDY 1

The proposed study will involve a very close replication of Moss-Racusin et al. (2012). The raw data from this study (once collected) will be available online (<https://osf.io/sefn8/>).

METHOD

All methods have received IRB approval from Rutgers University Institutional Review Board (IRB Approval #2018001059). The study is pre-registered on PsyArXiv via a posted pre-print of this registered replication report (<https://psyarxiv.com/ezp6d>), and the instrument will be hosted on a *Qualtrics* site. All measures and materials are available online (<https://osf.io/sefn8/>).

Criteria For Hypothesis Confirmation And Replication

It is important to have clear, a priori standards for determining whether or not something being studied replicates or not. But, this turns out to be more complex than it may seem at first. For example, a simple approach to replication would be, if an original finding is significant (using arbitrary thresholds), then a replication would be considered successful if it produced a statistically significant result in the same direction. This is problematic, though, as underpowered replications may fail to replicate possible true effects (Simonsohn, 2015). Because of complexities such as this, considerable effort has been made over the last several years to establish criteria to address and solve this and other problems. Since there is no universally accepted or recognized standard for what constitutes replication success or failure, in the present research we will use two different approaches for determining whether the original effects replicate: Small telescopes (Simonsohn, 2015), and an effect size consistency and precision approach (LeBel, Vanpaemel, Cheung, & Campbell, 2019). Standard frequentist statistics also will be used for this pre-registered study, as that is what was used in the original. But, given the large target sample size (discussed below), effect sizes will be the primary lens through which

the importance of any findings are determined using these two approaches, and by which we will classify whether the original effects replicate or not.

Small Telescopes Criteria

Simonsohn (2015) suggested that researchers evaluate whether results obtained are consistent with an effect size that would have been big enough to have been detected in the original study. Specifically, is the effect size obtained in the replication statistically smaller than an effect the original study had 33% power to detect? If so, Simonsohn (2015) details, the finding did not replicate because it “is inconsistent with the notion that the studied effect is large enough to have been detectable with the original sample size” (p. 562). If, however, the effect size obtained in the replication is not statistically smaller than an effect the original study had 33% power to detect, the replication would be considered successful.

Power analyses were conducted using G*Power. Following Simonsohn (2015), the effect sizes Moss-Racusin et al. (2012) had 33% power to detect were calculated for t-tests, difference between two independent means, one-tailed. Given $n_1=63$, $n_2=64$, Moss-Racusin et al. (2012) had 33% power to detect $d=.22$ for their dependent variables. Effect sizes for the DVs obtained by the replication will be evaluated against this effect size to determine if they are smaller, or not. If any 95% confidence interval surrounding the obtained effect sizes is below the small effect size of $d=.22$, that effect would be considered to have failed to replicate. This, then, would indicate that the original study could not have meaningfully examined and obtained the effect obtained in the replication. Effect sizes with 95% confidence intervals that include or are above $d=.22$, by this criteria, would be interpreted as replication successes.

LeBel et al. Criteria

LeBel et al. (2019) detail a second principled statistical approach for evaluating the degree of success a study has had in replicating a prior study. In their system, there are three questions that must be answered:

1. Does the replication study provide signal or noise? Signal/no signal refers to whether or not the 95% CI around the effect size obtained by the replication contains zero, which is equivalent to a traditional null hypothesis test using $p < .05$ as the threshold for significance. If so, there is no “signal.” If the CI does not contain zero, there is “signal.”
2. Are the findings of the replication consistent or inconsistent with the original? This is addressed by determining whether the 95% CI around the effect size obtained by the replication contains the original study effect size point estimate. If so, the results are consistent; otherwise, they are not.
3. If the replication’s results have signal and are consistent with the original, the conclusion is that the second study has replicated the findings of the first. If the replication’s results have signal and are *inconsistent* with the original, it next must be determined if the results are opposite or not. That is, it must be determined whether the replication effect size is on the same side of zero as the original. If the replication effect size is on the other side of zero, it is in the “wrong” direction and the second study did not replicate the first. If the replication effect size is on the same side of zero, then a final question should then be addressed: Is the replication effect size larger or smaller than the original?

Following LeBel et al. (2019), 95% confidence intervals will be calculated for the effect sizes of the dependent variables in the replication. These effect sizes and their 95% CIs will then be evaluated and clasified using the criteria detailed by LeBel et al. (2019) to determine the outcome of the replication.

Jointly Assessing the Two Criteria

In an ideal world, these two methods (Small Telescopes and the LeBel et al. criteria) will yield a clear and identical conclusion: either our project succeeded or failed to replicate Moss-Racusin et al. (2012). Statistical tests provide answers that are only probabilistic, not certain, and different types of tests can provide different answers. It is, therefore, possible that our study will not provide a clear answer to the question “Did it replicate?” The two methods may produce different answers. The LeBel et al. (2019) method, especially, is quite nuanced and can produce a wide range of answers, reflecting greater or lesser degrees of replication. If the results of the present study are high in ambiguity—they do not clearly indicate whether the replication succeeded or failed—that is exactly what we plan to report.

Criteria For Evidence Of Gender Bias

Regardless of our conclusions about replication, a separate issue is whether this project’s results will constitute evidence of gender bias. Inasmuch as the difference in ratings of male and female lab manager applicants is unlikely to be exactly zero, the issue can be viewed as requiring identification of the smallest effect size of interest (Lakens, 2017). How large is a nontrivial bias effect? Although there is no objective gold standard for addressing this question, prior research on both gender differences and discrimination has provided some useful and converging guidelines. Hyde (2005) advanced the “gender similarities hypothesis” on the grounds that most psychological gender differences were small, and she used a figure displaying the very large overlap between two distributions that differed at $d=.21$ to illustrate the similarity of the populations. We interpret this to mean ‘ $d<.21$ is consistent with the gender similarities hypothesis.’ This conclusion converges with that reached in research on discrimination, such as that, as argued by Greenwald, Banaji, and Nosek (2015), differences greater than or equal to an effect size of $r=.111$ are socially important and reflect significant discrimination.

For the purposes of the present research, we will use the cutoff of $d=.21$ (which is approximately equal to $r=.111$) for determining whether our research obtains evidence of gender bias. To evaluate this cut-off statistically, we will employ an inferiority test (Lakens, Scheel, & Isager, 2018). Specifically, using a one-sided t-test, we will compare the obtained effect sizes to our smallest effect sizes of interest (SESOI; $d=.21$). From this test there will be three possible outcomes.

1. 95% confidence intervals that at the lower bound, are greater than or equal to $d=.21$ would constitute clear evidence of gender bias.
2. 95% confidence intervals that are below and don't include $d=.21$ would constitute evidence inconsistent with gender bias.
3. 95% confidence intervals that include d above or equal to $.21$, but also values below $.21$, would provide ambiguous evidence of gender bias--neither clear evidence for, or clear evidence to the contrary.

Sample and Recruitment Strategy

Sample Size Determination: We are planning to recruit a sample of 564 participants, in comparison to the 127 participants recruited by Moss-Racusin et al. (2012). Our power analysis, using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007), for a one-sided t-test (since we are interested in effects at or above our specified effect size) gives us 99.99% power at a 5% error rate to detect the smallest effect size obtained by Moss-Racusin et al. (2012) (i.e., $d=.60$; see Table 2). Power to test our SESOI ($d=.21$) with 564 participants is 80.10% (i.e., meets conventional standards).

Table 2				
Description	Power	Alpha	Effect Size	N
Power for our proposed sample size, given smallest effect size obtained by Moss-Racusin et al. (2012)	99.99%	.05	$d=.60$	564
Power to test our SESOI with $n=564$	80.10%	.05	$d=.21$	564
<i>Note 1:</i> All power analyses conducted for one-sided t-tests since we are interested in effects at, or above, our specified effect size <i>Note 2:</i> $d=.60$ reflects the smallest effect size obtained by Moss-Racusin et al. (2012), $d=.21$ reflects our SESOI				

Following Moss-Racusin et al. (2012), we strategically sought to recruit faculty from universities that have departments representative of high-quality United States science programs. A geographically dispersed sample of 40 public and 15 private “R1” universities (i.e., doctoral institutions with the highest research activity), as classified by The Carnegie Classification of Institutions of Higher Education (n.d.), were selected for faculty recruitment.

Within each university, in our attempt to directly replicate Moss-Racusin et al. (2012)’s methods, participants will be recruited exclusively from biology, chemistry, and physics departments. For each selected science department, publicly available department web-pages and online rosters were used to create a list of eligible faculty participants. Participants had to meet the same a priori qualifications employed by Moss-Racusin et al. (2012); they had to be clearly identified as active tenure-track professors (i.e., full, associate, or assistant professor). Consequently, emeritus, visiting, adjunct, and research professors were not included, nor were lecturers, postdocs, or other instructional staff. Additionally, those with invalid email addresses, those who had no email addresses listed, and those who were identified as colleagues of the authors were also not included.

In total, a list of 5866 potential science faculty participants has been generated. Participant compensation was not described in the original publication or it’s supplementary

materials. But, to be consistent with procedures described by Moss-Racusin in private email correspondence (October 16, 2019), to incentivize participation each participant will be offered \$30 for completing the survey (inflation adjusted from the \$25 per participant used by Moss-Racusin). Following other recent studies of university faculty (e.g., Honeycutt & Freberg, 2017; Williams & Ceci, 2015), it is anticipated that approximately 7-30% of these professors will participate. As such, to reach the desired sample size as determined by our a priori power analysis, we will randomize our list of potential science faculty participants and begin by contacting 600 of them. From there we will contact participants in waves of 50 until we obtain completed surveys from at least 564 participants. The remaining balance of emails for potential science faculty participants (anticipated to be approximately 3866 emails) will be used in Study 2. Data collection will commence immediately upon this registered report successfully obtaining in principle acceptance (Stage 1 IPA). Timeline for completion of the study, and Stage 2 submission, is estimated to be four months after receiving Stage 1 IPA.

Design and Manipulation

Participants will be randomly assigned to one of two student applicant gender conditions: that of “John” or “Jennifer.” Both conditions are identical to those used by Moss-Racusin et al. (2012), with two difference: 1- dates have been updated to be current (i.e., 2011 updated to 2020); and 2- GRE scores have been converted to scores in the same percentiles, but now in the new scoring format.

Participants will see only one set of materials (male or female), and both profiles are exactly identical, except for the applicant’s gender—one using “John” and male pronouns, and the other using “Jennifer” and female pronouns. The student lab manager applications will be prefaced with the same introduction used in the original research, which was designed to bolster

the credibility and adjust for differences in expectations and practices regarding lab managers that may differ between academic areas.

Measures

All measures utilized in the present study will be identical to those used by Moss-Racusin et al. (2012) . The full set of materials are available online (<https://osf.io/sefn8/>). Additional suspicion check questions not employed in the original study will be used to aid in assessing the possible impact of prior knowledge about the original study (described below in detail).

Demographic Variables

Participants will provide demographic information that includes their gender (1=male, 2=female, 3=other/decline), age, ethnicity/race, political ideology (1=very liberal, 4=moderate, 7=very conservative), rank (assistant, associate, full, or distinguished professor), and academic area (biological sciences, chemistry, physics). For the purpose of being diligent in masking the identities of participants to the best of our ability, after conducting analyses, the age variable in the publicly posted datafile will be converted to categories as follows: 1-under 18, 2- 18 to 27, 3- 28 to 37, 4- 38 to 47, 5- 48 to 57, 6- 58 to 67, 7- 68 to 77, 8- 78 and older.

Applicant Competence

Competence of the applicants in the two profiles will be assessed using the same three items used by Moss-Racusin et al. (2012). Faculty participants will be asked: 1-Did the applicant strike you as competent? 2-How likely is it that the applicant has the necessary skills for this job? 3-How qualified do you think the applicant is? (1=not at all, 7=very much). Higher scores reflect a perception of higher competence.

Applicant Hireability

Using the same three items used by Moss-Racusin et al. (2012), faculty participants will rate how hireable they view the student applicant to be for a lab manager position. Items include:

1-How likely would you be to invite the applicant to interview for the laboratory manager job? 2-How likely would you be to hire the applicant for the laboratory manager job? 3-How likely do you think it is that the applicant was actually hired for the laboratory manager job he/she applied for? (1=not at all likely, 7=very likely). Higher scores reflect a perception of increased likelihood of hireability.

Mentoring

Three items, identical to those used by Moss-Racusin et al. (2012), will be used to assess faculty participants' willingness to mentor the applicant. Items include: If you encountered this student at your own institution, how likely would you be to... 1-Encourage the applicant to stay in the field if he/she was considering changing majors? 2-Encourage the applicant to continue to focus on research if he/she was considering switching focus to teaching? 3-Give the applicant extra help if he/she was having trouble mastering a difficult concept? (1=not at all likely, 7=very likely). Higher scores reflect a greater willingness to mentor the student applicant.

Salary Conferral

A single item, identical to that used by Moss-Racusin et al. (2012), will be used to assess the recommended salary for the student applicant. Faculty participants will be asked: If you had to choose one of the following starting salaries for the applicant, what would it be? Participants will be given the following response options: 1-(\$15,000), 2-(\$20,000) 3-(\$25,000), 4-(\$30,000), 5-(\$35,000), 6-(\$40,000), 7-(\$45,000), 8-(\$50,000).

Subtle Gender Bias

Moss-Racusin et al. (2012) included the Modern Sexism Scale to measure subtle bias against women (Swim et al., 1995). Higher scores can be loosely characterized as modern prejudice towards women, indifference towards women's rights, or subtle bias against women.

Likeability

Faculty participants will also be asked to indicate the extent to which they (dis)liked the student applicant. The three items, identical to those used by Moss-Racusin et al. (2012) include: 1-How much did you like the applicant? 2-Would you characterize the applicant as someone you want to get to know better? 3-Would the applicant fit in well with other laboratory members? (1=not at all, 7=very much). Higher scores reflect a greater liking of the student applicant.

Suspicion Checks

In an attempt to assess possible effects resulting from participants being aware of audit-type study methodologies or the original study, participants will be probed for suspicion. Using the same questions utilized by Moss-Racusin et al. (2012), participants will be asked what they believe was the purpose of the study, if they noticed anything unusual about the application they read, and if they felt deceived in any way. Responses will be rated for suspicion by three independent raters (1=not at all suspicious to 5=very suspicious). Ratings from the three independent raters (blind to condition) will be averaged together for a suspicion score for each participant. Participants who report suspicion about the study design (e.g., that this is an audit-type study), measures, awareness of the original study, or possible intent or objective of the study will not be excluded. Instead, as described later in planned analyses, tests will be conducted to assess if more suspicious participants respond differently than less suspicious ones.

Next, participants will be probed for cross-talk, using a question not used by Moss-Racusin et al. (2012). Specifically, participants will be asked if, prior to participating, they heard about or talked about this study with anyone else (yes/no). As described later in planned analyses, tests will be conducted to assess if those who indicate having engaged in cross-talk responded differently from those who did not on the main measures.

Additionally, questions will be incorporated at the end of the study (to ensure that everything prior is a direct replication) to assess whether participants are familiar with the Moss-Racusin et al. (2012) article. Participants will be presented with the complete bibliographic information of six research articles in randomized order, and asked to indicate how much they know about each article (*1=nothing at all, 5=I know a great deal about this*). This online supplementary material is available here: <https://osf.io/sefn8/>. One of the six articles will be Moss-Racusin et al. (2012), while the remaining five will be fakes in the sense that they do not actually exist; we will have concocted them. The fakes will be used to assess overclaiming (Paulhus, Harms, Bruce, & Lysy, 2003). Overclaiming refers to people claiming to know things they cannot possibly know because they do not exist.

We plan to assess whether overclaiming also characterizes respondents' claims about knowing about the Moss-Racusin et al. (2012) paper. For example, if about as many respondents claim to know of nonexistent studies as claim to know the Moss-Racusin et al. (2012) study, this would indicate random guessing rather than awareness. If more respondents indicate having heard of the Moss-Racusin et al. (2012) study than other studies, that difference would constitute one index of *how many* respondents actually had some familiarity with the study. For example, let's say 20% of the sample claims to have recognized the Moss-Racusin et al. (2012) study, and, on average, 15% claim to have recognized each of the other five fake studies. This says there is an ambient level of about 15% of participants saying, "Yes, I know that study" even when they cannot possibly know the study because it does not exist. If 15% is the ambient level of overclaiming, and if 20% stated that they recognized the Moss-Racusin et al. (2012) study, it would indicate that about 5% of the same actually did have some knowledge of the original study. In fact, this would be a low bar because even a vague memory of having heard of the study, without much recall of the actual findings, could be enough to trigger a "yes, I know that

study” response. Thus, this difference can be viewed as an *upper bound* on the likely extent to which awareness of the original study could have affected our results.

Manipulation Check

To assess how closely participants were paying attention, at the end of the survey, after overclaiming and suspicion check questions, participants will be asked one manipulation check questions: 1-What was the applicant’s gender? (1=male, 2=female, 3=not specified). This question will be mixed in with the following questions to avoid raising suspicion, and to disguise the primary question of interest: 1-How did you feel about the quantity of information available for you to evaluate the applicant? (1-7, 1=far too little, 4=neither too little nor too much, 7=far too much). 2-What was the applicant’s ethnicity/race? 3-Was there any information missing from the application summary that would otherwise help you evaluate the applicant? 4-What did you perceive the applicant’s personal political beliefs/orientation to be? (1-7, 1=very liberal, 4=moderate/neutral, 7=very conservative, plus 8=don’t know—not enough information, or unable to infer).

Procedure

The first author will email participants. The email will contain the cover story describing why we are soliciting their assistance, along with a link to the online materials. To reduce contamination and the possibility that participants will discuss the survey with others, participants will have two weeks to complete the study online at their convenience. After consenting to participate, participants will first read the introduction. Following, participants will be randomly assigned to review one of the two student profiles: that of John or Jennifer. They will then evaluate the profile by completing the aforementioned measures for competence, hireability, mentoring, salary conferral, subtle gender bias, and likeability. Last, participants will complete demographic measures, the funneled debriefing, the suspicion check questions, and

then will be debriefed. The first author will send participants a reminder email after one week has elapsed, and all surveys will close two weeks after the survey was sent.

PLANNED DATA ANALYSIS

Our planned data analysis will consist of three phases. Phase one will consist of preliminary analyses such as frequencies and creating composite variables from the scales. Phase two will consist of direct replication analyses. Phase three will consist of additional interesting and useful analyses not conducted or not reported by Moss-Racusin et al. (2012).

Phase 1: Preliminary Analyses

Frequencies will be reported for demographics. All scale creation methods follow Moss-Racusin et al. (2012)'s procedures exactly. Items comprising scales will be averaged. Scale scores will be created as follows: Eight composite variables will be constructed from: 1-the four applicant competence questions; 2-the three questions related to applicant hireability; 3-the three questions related to mentoring; 4-the eight modern sexism scale items for subtle gender bias; 5-the three questions related to likeability; 6-the standardized salary variable will be averaged with the four applicant competence items to form a generalized competence index item; 7-the six questions asking how much participants know about various research articles (one real, five fake, unknown to participants) will be used to form a variable for overclaiming. Following Paulhus, Harms, Bruce, and Lysy (2003), the variable for overclaiming will be calculated as follows: First, we will calculate the proportion of real articles ("P-Real") that were given a rating higher than one (one here is our cutoff for proportion of "hits"). Next, we will calculate the proportion of fake articles ("P-FA") given a rating higher than one (one here is our cutoff for proportion of "false alarms" or "foils"). Finally, overclaiming (i.e., "response bias") will be calculated by adding P-Real and P-FA ; 8-the three ratings of suspicion for each participant from the three

independent raters. Descriptives (means, SDs, ranges) and correlations will be reported for all scales and non-categorical variables.

Phase 2: Direct Replication Analyses

Gender Bias

Our first set of analyses will examine whether there is gender bias in perceptions of competence, hireability, and mentoring, in addition to salary conferral. Moss-Racusin et al., (2012, p. 16475) assessed this as follows:

The competence, hireability, salary conferral, and mentoring scales were each submitted to a two (student gender; male, female) $\times$ two (faculty gender; male, female) between-subjects ANOVA. In each case, the effect of student gender was significant (all $P < 0.01$), whereas the effect of faculty participant gender and their interaction was not (all $P > 0.19$).

Therefore, to follow their approach exactly, the composite variables for competence, hireability, and mentoring, in addition to the single item for salary conferral, will each be submitted to separate 2 (applicant gender: male, female) $\times$ 2 (faculty participant gender: male, female) between subjects ANOVAs.

To follow this up, Moss-Racusin et al., (2012, p. 16475) then conducted main effects for competence, hireability, mentoring, and salary conferral. In the following excerpt, these are described as “simple effects,” but in email correspondence with Moss-Racusin (March 14, 2019), it was clarified that these were actually “main effect tests,” which were mis-reported as “tests of simple effects”:

Tests of simple effects (all $d > 0.60$) indicated that faculty participants viewed the female student as less competent [$t(125) = 3.89$, $P < 0.001$] and less hireable [$t(125) = 4.22$, $P < 0.001$] than the identical male student (Fig. 1 and Table 1). Faculty participants also

offered less career mentoring to the female student than to the male student [$t(125) = 3.77, P < 0.001$]. The mean starting salary offered the female student, \$26,507.94, was significantly lower than that of \$30,238.10 to the male student [$t(124) = 3.42, P < 0.01$]. Therefore, main effects tests (i.e., four independent-samples t-tests) will be conducted to compare competence, hireability, mentoring, and salary conferral, related to the male versus the female applicant profile. Though main effects will be available in the output from the 2 x 2 ANOVAs described earlier, these t-tests will be conducted so as to closely follow what was conducted and reported in the original, as described in the quoted excerpt. That said, we will also report both main effects and the interaction from the ANOVAs.

Moss-Racusin et al., (2012, p. 16475-16476) also describe a series of analyses in such a manner that it is difficult to recreate what they actually did. Here is the full quote:

In support of hypothesis B, faculty gender did not affect bias (Table 1). Tests of simple effects (all $d < 0.33$) indicated that female faculty participants did not rate the female student as more competent [$t(62) = 0.06, P = 0.95$] or hireable [$t(62) = 0.41, P = 0.69$] than did male faculty. Female faculty also did not offer more mentoring [$t(62) = 0.29, P = 0.77$] or a higher salary [$t(61) = 1.14, P = 0.26$] to the female student than did their male colleagues.

Table 1 clearly reports bias effects--differences in evaluations of male versus female applicants for competence, hireability, and mentoring. However, Moss-Racusin et al.'s next statement, about simple effects, *only* refers to gender of faculty differences in ratings of female applicants. What does this mean? They had (previously) reported that the gender of faculty by gender of applicant interaction had a $p > .19$. But what about the simple effects? Those are only reported for faculty gender differences in evaluations of the female applicant. What is left unmentioned is

whether they also performed a similar analysis for evaluations of the male applicant (which we will conduct in Phase 3, described below).

Therefore, for the direct replication portion of this study to tightly conform to the analyses reported by Moss-Racusin et al. (2012), four additional simple effects tests (i.e., independent samples t-tests) will be conducted to compare male versus female faculty respondent ratings of female applicant competence, hireability, mentoring, and salary conferral.

Moss-Racusin et al., (2012, p. 16476) also sought to determine if faculty scientific field, age, or tenure status affected bias:

In addition, faculty participants' scientific field, age, and tenure status had no effect (all $P > 0.53$). Thus, the bias appears pervasive among faculty and is not limited to a certain demographic subgroup.

It is unclear what analyses Moss-Racusin et al. (2012) conducted, and whether the results reported in the quote refer to main effects, interactions, or both. Therefore, four two-way ANOVAS (applicant gender by scientific field/age/tenure status) will be conducted to determine if there are differences by scientific field, age, or tenure status in ratings of the female profile and in ratings of the male profile on competence, hireability, mentoring, and salary conferral. Both main effects and the interaction will be reported, which will ensure that we will have conducted every analysis Moss-Racusin et al. (2012) conducted.

Last, Moss-Racusin et al. (2012) assessed if there was a significant difference in ratings of likability for the male versus female applicant, described as follows (Moss-Racusin et al., 2012, p. 16477):

Faculty participants reported liking the female (mean = 4.35, SD = 0.93) more than the male student [(mean = 3.91, SD = 0.108), $t(125) = -2.44$, $P < 0.05$].

This difference is the main effect from the 2x2 ANOVA conducted earlier in this section earlier, so we will report that result. However, Moss-Racusin et al. (2012) reported conducting a t-test rather than the main effect of the ANOVA. Therefore, we will also analyze as a t-test, such that a main effect test (i.e., one independent-samples t-test) will be conducted to compare the likeability of the male versus the female applicant profile.

Mediation and moderation analyses

Moss-Racusin et al. (2012) conducted a series of mediation and moderation analyses to determine if hiring discrimination against the female applicant would be mediated by perceptions that the female applicant is less competent than the identical male applicant, or by likeability.

The analyses were described as follows (Moss-Racusin et al., 2012, p. 16476 & 16477):

The initially significant impact of student gender on hireability ($\beta = -0.35, P < 0.001$) was reduced in magnitude and dropped to nonsignificance ($\beta = -0.10, P = 0.13$) after accounting for the impact of student composite competence (which was a strong predictor, $\beta = 0.69, P < 0.001$), Sobel's $Z = 3.94, P < 0.001$ (Fig. 3). This pattern of results provides evidence for full mediation, indicating that the female student was less likely to be hired than the identical male because she was viewed as less competent overall... Only composite competence (and not likeability) helped to explain why the female student was less likely to be hired; in mediation analyses, student gender condition ($\beta = -0.48, P < 0.001$) remained a strong predictor of hireability along with likeability ($\beta = 0.60, P < 0.001$).

Therefore, to follow their approach, a series of mediation analyses will be conducted to test the proposed processes. First, to examine the mediating roles of perceived competence and likeability, we will conduct a parallel mediation analysis using PROCESS for SPSS v3.3 Model 4. In this analysis, applicant gender will be the independent variable, hireability the dependent

variable, and competence and likeability will be the mediators. To replicate their results, we would need to find that composite competence strongly correlates with hireability and controlling for it reduces the effect of applicant gender on hireability to near 0.

The mediation analyses in the Moss-Racusin et al. (2012) study were followed by moderation analyses to determine if pre-existing biases held by faculty (assessed using the Modern Sexism Scale) would predict ratings of the male and female applicants. The analyses were described as follows (Moss-Racusin et al., 2012, p. 16476):

Results of multiple regression analyses indicated that participants' preexisting subtle bias against women significantly interacted with student gender to predict perceptions of student composite competence ($\beta = -0.39, P < 0.01$), hireability ($\beta = -0.31, P < 0.05$), and mentoring ($\beta = -0.55, P < 0.001$). To interpret these significant interactions, we examined the simple effects separately by student gender. Results revealed that the more preexisting subtle bias participants exhibited against women, the less composite competence ($\beta = -0.36, P < 0.01$) and hireability ($\beta = -0.39, P < 0.01$) they perceived in the female student, and the less mentoring ($\beta = -0.53, P < 0.001$) they were willing to offer her. In contrast, faculty participants' levels of preexisting subtle bias against women were unrelated to the perceptions of the male student's composite competence ($\beta = 0.16, P = 0.22$) and hireability ($\beta = 0.07, P = 0.59$), and the amount of mentoring ($\beta = 0.22, P = 0.09$) they were willing to offer him.

Therefore, to test for the predicted moderations, a series of multiple regression tests will be conducted to test whether the association between applicant gender and competence (using the generalized competence index term), hireability, and mentoring differs depending on participants' responses on the Modern Sexism Scale. Interactions with negative slopes (i.e., negative β 's) will confirm that the relationship between applicant gender and competence,

hireability, and mentoring differs as a function of responses to the MSS (i.e., subtle bias toward women). To interpret interactions, simple effects tests will be conducted separately by applicant gender. Additionally, though not conducted by Moss-Racusin et al. (2012), interactions will be further analyzed by computing simple effects at ± 1 standard deviation of the moderator, and a visual graph plotting the slopes will be generated.

Phase 3: Additional Analyses

Gender differences—ratings of male applicant

Earlier, it was described how Moss-Racusin et al. (2012) conducted a series of simple effects tests for faculty ratings by gender of the female applicant. Missing, though, is the reporting of tests for gender differences among faculty in ratings of the male applicant. Moss-Racusin et al. (2012, Table 1) reported means for the faculty ratings by gender of male applicants, but the simple effects and additional results were not reported. Therefore, four simple effects tests (i.e., independent samples t-tests, one for each dependent variable) will be conducted to compare male versus female faculty respondent ratings of male profile competence, hireability, mentoring, and salary conferral.

Tests beyond those in Moss-Racusin et al. (2012)—overclaiming, familiarity, and suspicion

A series of progressively more aggressive tests will be conducted to assess whether familiarity with the original study may have influenced responses in the present study. As noted earlier, participants will answer one question assessing how familiar (if at all) they are with Moss-Racusin et al. (2012) and five questions assessing overclaiming—the extent to which they claim they are familiar with nonexistent studies (see, e.g., Paulhus et al., 2003). If participants claim to be familiar with Moss-Racusin et al. (2012) at about the same rate as with nonexistent studies (tested with a one-sample chi-square), we will conclude that, though some participants

may overclaim, there is no evidence that, in general, our respondents actually *are* familiar with Moss-Racusin et al. (2012), despite the attention it received.

To test for the possible impact of overclaiming, multiple regressions will examine whether the amount of gender bias in judgements (i.e., competence, hireability, mentoring, and salary conferral, each in a separate test) differs depending on overclaiming. A possible hypothesis here is that overclaiming might reflect some sort of social desirability bias or motivation to appear unsexist, such that more overclaiming would predict less bias. Of course, at least hypothetically, there are many other possibilities, including overclaiming predicts reverse bias (more overclaiming, more bias favoring women) or even an entirely opposite pattern altogether (more overclaiming predicts more bias favoring men). We do not have an a priori prediction. Thus, the exploratory question here is: does how much subjects overclaim moderate findings? Interaction effects between overclaiming and target gender in the models would confirm that the relationship between the application rated and each DV differs as a function of overclaiming. Simple effects will be computed at ± 1 standard deviation of the moderator and a visual graph plotting the slopes will be generated.

If more people claim to be familiar with Moss-Racusin et al. (2012) than with nonexistent studies, additional steps will be taken to examine how this might be related to responses. Specifically, a series of four regressions will be conducted to examine whether gender bias depends on the level of familiarity with Moss-Racusin et al. (2012). Specifically, a series of moderated multiple regression models will be tested to determine whether the amount of sex bias in judgements (i.e., competence, hireability, mentoring, and salary conferral, each in a separate test) differs depending on familiarity with Moss-Racusin et al. (2012). Interaction effects between familiarity and target gender would indicate that the relationship between the application rated and each DV differs as a function of familiarity with Moss-Racusin et al.

(2012). Simple effects will be computed at ± 1 standard deviation of the moderator and a visual graph plotting the slopes will be generated. If responses are affected by familiarity with the original research study, we would expect familiarity with the original research study to moderate any obtained effects, assuming familiarity affects judgments in ways that lead to socially desirable responding (but, this concern may be overrated. See, e.g., Zwaan et al., 2018).

Additional exploratory analyses will examine whether level of suspicion predicted participant responses. Multiple regression models will be conducted to determine whether the amount of gender bias in judgements (i.e., competence, hireability, mentoring, and salary conferral, each in a separate test) differs depending on the level of suspicion. That is, does participant suspicion (see earlier suspicion check section) moderate findings? Simple effects will be computed at ± 1 standard deviation of the moderator and a visual graph plotting the slopes will be generated. Possible outcomes for each of the four DVs could include: 1-suspicion not moderating the effects (i.e., suspicion not impacting our tests for gender bias); 2-suspicion moderating the effects, such that gender bias effects are weaker for more suspicious participants; or 3-suspicion moderating the effects, such that gender bias effects are stronger for more suspicious participants.

Finally, an additional exploratory analysis will examine whether cross-talk (participants who have versus have not acknowledged talking about the present study with others) predicted bias. Specifically, we will conduct a 2 (applicant gender: male, female) $\times$ 2 (cross talk: yes, no) between subjects ANOVAs for competence, hireability, mentoring, and salary conferral. A main effect of cross-talk, or an interaction, would indicate that participants who talked to others about the study somehow responded differently from those who did not. In sum, our a priori tests for overclaiming, familiarity, suspicion, and cross talk provide a robust framework for assessing

the role of social desirability and related biases--if any--in our assessment of evaluations of male and female lab manager applicants.

STUDY 2

This second proposed study will involve an extension of Moss-Racusin et al. (2012). This larger study will expand the sample to include faculty in disciplines well beyond those included in the original study. Specifically, despite the fact that the Moss-Racusin et al. (2012) study is often cited as showing gender bias in STEM, it actually only showed gender bias in S(cience) because it did not assess bias among faculty in “TEM.” Because Study 1 will be a very close replication of Moss-Racusin et al. (2012), it, too, will have only assessed gender bias in S. Study 2 addresses this limitation by performing a very close replication of Study 1 (and Moss-Racusin et al., 2012) with one crucial difference: our target participant sample will be STEM faculty broadly, not just S(cience) faculty. Thus, even though an extension, it will still be a very close replication of the original,. Unlike Study 1, and due limited funding and ambiguity about whether compensation affects responses (e.g., it wasn’t described in the original study or its supplementary materials) participants will not be compensated for participation. The raw data from this study (once collected) will be available online (<https://osf.io/sefn8/>).

METHOD

All methods have received IRB approval from the Rutgers University Institutional Review Board (IRB Approval #2018001059). The study will be pre-registered on Open Science Framework, and the instrument will be hosted on a *Qualtrics* site. All measures and materials are available online (<https://osf.io/sefn8/>).

The design, manipulations, procedures, measures, and criteria (e.g., for hypothesis confirmation, replication, and gender bias) will be identical those of Study 1 with the exception of sample recruitment. The full set of materials are available online (<https://osf.io/sefn8/>).

Sample and Recruitment Strategy

Sample Size Determination: We anticipate recruiting at least 1075 “STEM” participants, which would constitute a 7% response rate. Our power analysis, using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007), for a one-sided t-test (since we are interested in effects at or above our SESOI) gives us 99.99% power at a 5% error rate to detect the smallest effect size obtained by Moss-Racusin et al. (2012) (i.e., $d=.60$; see Table 3). Power to test our SESOI ($d=.21$) with 1075 participants is 96.37% (i.e., exceeds conventional standards).

Table 3				
Description	Power	Alpha	Effect Size	N
Power for our proposed “STEM” sample size, given smallest effect size obtained by Moss-Racusin et al. (2012)	99.99%	.05	$d=.60$	1075
Power to test our SESOI for “STEM”	96.37%	.05	$d=.21$	1075
<i>Note 1:</i> All power analyses conducted for one-sided t-tests since we are interested in effects at, or above, our specified effect size <i>Note 2:</i> $d=.60$ reflects the smallest effect size obtained by Moss-Racusin et al. (2012), $d=.21$ reflects our SESOI				

In Study 1, we drew from the pool of 5866 science faculty emails (i.e., email addresses for biology, chemistry, and physics faculty) to recruit the desired sample of 564 participants. Inasmuch as biology, chemistry, and physics departments are science fields, but not technology, engineering, or mathematics, we refer to them as the “S” portion of STEM. Our extension here will involve recruiting “TEM” participants in addition to “S” participants. To capture “TEM”, faculty participants will be recruited from engineering, computer science, mathematics, and statistics departments. These additional emails have been collected from the same geographically dispersed sample of 40 public and 15 private “R1” universities utilized in Study 1, using the same collection and inclusion criteria specified in Study 1. Thus, this extends recruitment

compared to Moss-Racusin et al. (2012), as they only recruited from science programs (i.e., biology, chemistry, and physics departments).

In total, a list of 11,496 “TEM” email addresses will comprise the list of potential “TEM” faculty participants for Study 2. These will be added to the balance of unused (i.e., uncontacted) emails from Study 1 [Given an anticipated response rate of 30% in Study 1, we expect to contact about 2000 of the 5866 science faculty (“S”) to obtain 564 responses. This, then, leaves an anticipated balance of about 3866 uncontacted science faculty]. In total, we anticipate contacting approximately 15,362 STEM faculty. Following other recent studies of university faculty (e.g., Honeycutt & Freberg, 2017; Williams & Ceci, 2015), it is anticipated that approximately 7-30% of these professors will participate.

Data collection will commence immediately upon this registered report successfully obtaining in principle acceptance (Stage 1 IPA), and upon the completion of data collection for Study 1. Timeline for completion of the study, and Stage 2 submission, is estimated to be four months after receiving Stage 1 IPA.

PLANNED DATA ANALYSIS

Our planned data analysis will consist of three phases. Phase 1 will consist of preliminary analyses such as frequencies and creating composite variables from the scales. Phase 2 will consist of direct replication analyses. In this phase an exact repeat of Moss-Racusin et al. (2012)’s analyses will be conducted on this “STEM” sample of faculty. Phase 3 will consist of additional interesting and useful analyses not conducted/reported by Moss-Racusin et al. (2012), conducted here on “STEM” faculty participants.

Phase 1: Preliminary Analyses

As in Study 1, frequencies will be reported for demographics. Scale creation will be identical to Study 1. Descriptives (means, SDs, ranges) and correlations will be reported for all scales and non-categorical variables.

Phase 2: Very Close Replication Analyses—STEM

Phase 2 analyses will be conducted exactly as conducted in Study 1 Phase 2, using the present sample of STEM participants.

Phase 3: Additional Analyses

Phase 3 analyses will be conducted exactly as conducted in Study 1 Phase 3, using the present sample of STEM participants. Additionally, and unique to Study 2, we will conduct a series of 2 (gender of student applicant: male, female) x 2 (faculty participant grouping: “S”, “TEM”) between-subjects ANOVAs for competence, hireability, and mentoring, in addition to the single item for salary conferral, to assess possible differences between “S” and “TEM.”.

STUDY 3

Study 3 will be a robust meta-analysis. Because Studies 1 and 2 are nearly identical, and administer identical measures, it will also be a clean meta-analysis. That is, Study 3 eliminates the “combining apples and oranges” problem of combining studies with very different outcomes. Furthermore, there will be no need to combine effect sizes. Instead, we can directly perform the main analyses. Simply put, this third study will combine all data from Study 1 and Study 2. It will then use conventional statistics to examine, across all samples, if there is evidence of gender bias in STEM fields at R1 universities in the United States.

METHOD

Data for this study will consist of the pooled data from the prior two studies in this paper.

Power

We anticipate the full data set from Studies 1 and 2 to consist of 1639 STEM faculty participants (564 “S” Study 1 + 1075 “STEM” Study 2). Our power analysis, using G*Power (Faul, Erdfelder, Lang, & Buchner, 2007), for a one-sided t-test (since we are interested in effects at or above our SESOI) gives us 99.99% power at a 1% error rate to detect the smallest effect size obtained by Moss-Racusin et al. (2012) (i.e., $d=.60$; see Table 4). Power to test our SESOI ($d=.21$) with 1637 participants is 99.54% (i.e., exceeds conventional standards).

Table 4				
Description	Power	Alpha	Effect Size	N
Power for our expected sample size, given smallest effect size obtained by Moss-Racusin et al. (2012)	99.99%	.05	$d=.60$	1639
Power to test our SESOI	99.54%	.05	$d=.21$	1639
<i>Note 1:</i> All power analyses conducted for one-sided t-tests since we are interested in effects at, or above, our specified effect size <i>Note 2:</i> $d=.60$ reflects the smallest effect size obtained by Moss-Racusin et al. (2012), $d=.21$ reflects our SESOI				

This total combined sample will consist of tenured or tenure track faculty from biology, chemistry, physics, engineering, computer science, mathematics, and statistics departments from the geographically dispersed sample of 40 public and 15 private “R1” universities (described in the prior studies).

Data analysis will commence immediately upon this registered report successfully obtaining in principle acceptance (Stage 1 IPA), and upon the completion of data collection for the prior three studies.

Very Close Replication Analyses for Gender Bias

The main analyses in Study 3 will provide an overarching test of whether our results reflect gender bias in STEM. Specifically, composite variables for competence, hireability, and mentoring, in addition to the single item for salary conferral, will each be submitted to separate 2 (applicant gender: male, female) x 2 (faculty participant gender: male, female) between subjects ANOVAs. Following, main effects tests (i.e., five independent-samples t-tests) will be conducted to compare competence, hireability, mentoring, salary conferral, and likeability ratings related to the male versus the female applicant profile. The gender bias hypothesis, as always, is that the male applicant will be evaluated more positively than the female applicant.

Finally, and unique to Study 3, we will conduct separate 2 (gender of student applicant: male, female) x2 [sample/study: Study 1 (“S”), Study 2 (“STEM”)] between-subjects ANOVAs for competence, hireability, and mentoring, in addition to the single item for salary conferral, to assess possible heterogeneity between the samples.

CONCLUSION

Gender gaps in STEM fields are widely considered to be an important social problem. One potential source of those gaps is gender bias among faculty. The proposed project includes two separate experimental tests for such biases and a meta-analysis. The first a direct replication of a classic demonstration of gender bias against women in academic science disciplines (Moss-Racusin et al., 2012), powered at a far higher level than the original (n of 564, compared to 127 in the original study), yielding over 99% power to detect the original effect size at $p < .05$ (and also over 99% power if one were to use the stringent standard of $p < .01$).

The second also includes a direct extension of that study to include technology, engineering and math (“TEM”) disciplines that were not included in the original study, in addition to the original science (“S”) disciplines. Given that Moss-Racusin et al. (2012) is cited

as evidence of bias in STEM and not just S(cience), it is clear that many researchers are very interested in reaching conclusions about gender bias in STEM and not just “S.” This research, therefore, can help fill the empirical gap in evidence regarding “TEM” left by the original study. This is important not only for replication purposes, but also for generalization: will Moss-Racusin et al.’s (2012) findings generalize to “TEM” fields, and STEM as a whole? Or will bias be larger, smaller, nonexistent, or even in the opposite direction when “TEM” fields are included? The only way to find out is to conduct the replication including “TEM” fields. This extension is also far more highly powered than the original study, also with over a 99% chance to find the original effect at $p < .05$ (and even at $p < .01$). Thus, this research is well positioned to provide the direct replication and generalization that is necessary to determine whether the effects obtained in an original small study are sufficiently robust that they can be readily detected in new research.

Of course, even if we fail to obtain the effects found in the original study, it does not mean anything is wrong with the original study. Random variation ensures that replications will rarely produce the exact same effect as an original study. Furthermore, society and culture change over time. Any differences between these studies and the original could reflect such societal changes rather than anything wrong with either the present studies or the original. Regardless, one cannot begin to address the question “have societal changes produced the different results?” until after it is known whether different results are actually obtained. Therefore, the present studies will provide an important contribution to understanding both the replicability of the original study and the current state of gender bias in STEM no matter what is found.

References

- Albers, C. J. (2015). Dutch research funding, gender bias, and Simpson's paradox. *Proceedings of the National Academy of Sciences*, 112(50), E6828–E6829.
<https://doi.org/10.1073/pnas.1518936112>
- Asendorpf, J. B., Conner, M., De Fruyt, F., De Houwer, J., Denissen, J. J. A., Fiedler, K., Fiedler, S., Funder, D. C., Kliegl, R., Nosek, B. A., Perugini, M., Roberts, B. W., Schmitt, M., Vanaken, M. A. G., Weber, H., & Wicherts, J. M. (2013). Recommendations for increasing replicability in psychology. *European Journal of Personality*, 27(2), 108–119. <https://doi.org/10.1002/per.1919>
- Beaudry, C., & Larivière, V. (2016). Which gender gap? Factors affecting researchers' scientific impact in science and medicine. *Research Policy*, 45(9), 1790–1817.
<https://doi.org/10.1016/j.respol.2016.05.009>
- Brown, A. J., & Goh, J. X. (2016). Some Evidence for a Gender Gap in Personality and Social Psychology. *Social Psychological and Personality Science*, 7(5), 437–443.
<https://doi.org/10.1177/1948550616644297>
- Budden, A. E., Tregenza, T., Aarssen, L. W., Koricheva, J., Leimu, R., & Lortie, C. J. (2008). Double-blind review favours increased representation of female authors. *Trends in Ecology & Evolution*, 23(1), 4–6. <https://doi.org/10.1016/j.tree.2007.07.008>
- Byars-Winston, A. (2014). Toward a Framework for Multicultural STEM-Focused Career Interventions. *The Career Development Quarterly*, 62(4), 340–357.
<https://doi.org/10.1002/j.2161-0045.2014.00087.x>
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T.-H., Huber, J., Johannesson, M., ... Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science

- between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637–644.
<https://doi.org/10.1038/s41562-018-0399-z>
- Carter, T. E., Smith, T. E., & Osteen, P. J. (2017). Gender comparisons of social work faculty using H-Index scores. *Scientometrics*, 111(3), 1547–1557.
<https://doi.org/10.1007/s11192-017-2287-0>
- Ceci, S. J., Ginther, D. K., Kahn, S., & Williams, W. M. (2014). Women in Academic Science: A Changing Landscape. *Psychological Science in the Public Interest*, 15(3), 75–141.
<https://doi.org/10.1177/1529100614541236>
- Ceci, S. J., & Williams, W. M. (2010). Sex differences in math-intensive fields. *Current Directions in Psychological Science*, 19(5), 275–279.
<https://doi.org/10.1177/0963721410383241>
- Ceci, S. J., & Williams, W. M. (2011). Understanding current causes of women's underrepresentation in science. *Proceedings of the National Academy of Sciences*, 108(8), 3157–3162. <https://doi.org/10.1073/pnas.1014871108>
- Chang, K. (2012, September 24). Bias Persists Against Women of Science, a Study Says. *The New York Times*. <https://www.nytimes.com/2012/09/25/science/bias-persists-against-women-of-science-a-study-says.html>
- Cohen, J. (1988). *Statistical Power Analysis for the Behavioral Sciences* (3rd ed.). Hillsdale, NJ: L. Erlbaum.
- Dion, M. L., Sumner, J. L., & Mitchell, S. M. (2018). Gendered Citation Patterns across Political Science and Social Science Methodology Fields. *Political Analysis*, 26(3), 312–327.
<https://doi.org/10.1017/pan.2018.12>
- Eagly, A. H., & Miller, D. I. (2016). Scientific Eminence: Where Are the Women? *Perspectives on Psychological Science*, 11(6), 899–904. <https://doi.org/10.1177/17456916166663918>

- Eaton, A. A., Saunders, J. F., Jacobson, R. K., & West, K. (2020). How gender and race stereotypes impact the advancement of scholars in STEM: Professors' biased evaluations of physics and biology post-doctoral candidates. *Sex Roles*, 82, 127–141.
<https://doi.org/10.1007/s11199-019-01052-w>
- Faul, F., Erdfelder, E., Lang, A.-G., & Buchner, A. (2007). G*Power 3: A flexible statistical power analysis program for the social, behavioral, and biomedical sciences. *Behavior Research Methods*, 39(2), 175–191. <https://doi.org/10.3758/BF03193146>
- Forscher, P. S., Cox, W. T. L., Brauer, M., & Devine, P. G. (2018). Little race or gender bias in an experiment of initial review of NIH R01 grant proposals.
<https://doi.org/10.31234/osf.io/r2xvb>
- Forscher, P. S., Lai, C. K., Axt, J. R., Ebersole, C. R., Herman, M., Devine, P. G., & Nosek, B. A. (2019). A meta-analysis of procedures to change implicit measures. *Journal of Personality and Social Psychology*, 117(3), 522–559.
<https://doi.org/10.1037/pspa0000160>
- Fraley, R. C., & Vazire, S. (2014). The N-Pact Factor: Evaluating the Quality of Empirical Journals with Respect to Sample Size and Statistical Power. *PLOS ONE*, 9(10), e109019.
<https://doi.org/10.1371/journal.pone.0109019>
- Funk, C., & Parker, K. (2018). *Women and Men in STEM Often at Odds Over Workplace Equity*. Pew Research Center. <http://www.pewsocialtrends.org/2018/01/09/women-and-men-in-stem-often-at-odds-over-workplace-equity/>
- Geraci, L., Balsis, S., & Busch, A. J. B. (2015). Gender and the h index in psychology. *Scientometrics*, 105(3), 2023–2034. <https://doi.org/10.1007/s11192-015-1757-5>

- Greenwald, A. G., Banaji, M. R., & Nosek, B. A. (2015). Statistically small effects of the Implicit Association Test can have societally large effects. *Journal of Personality and Social Psychology*, 108(4), 553–561.
- Handelsman, J., Cantor, N., Carnes, M., Denton, D., Fine, E., Grosz, B., Hinshaw, V., Marrett, C., Rosser, S., Shalala, D., & Sheridan, J. (2005). More Women in Science. *Science*, 309(5738), 1190–1191. <https://doi.org/10.1126/science.1113252>
- Honeycutt, N., & Freberg, L. (2017). The liberal and conservative experience across academic disciplines: An extension of Inbar and Lammers. *Social Psychological and Personality Science*, 8(2), 115–123. <https://doi.org/10.1177/1948550616667617>
- Hyde, J. S. (2005). The gender similarities hypothesis. *American Psychologist*, 60(6), 581–592.
- Ioannidis, J. P. A. (2008). Why Most Discovered True Associations Are Inflated. *Epidemiology*, 19(5), 640. <https://doi.org/10.1097/EDE.0b013e31818131e7>
- Kaatz, A., Gutierrez, B., & Carnes, M. (2014). Threats to objectivity in peer review: the case of gender. *Trends in Pharmacological Sciences*, 35(8), 371–373. <https://doi.org/10.1016/j.tips.2014.06.005>
- Knobloch-Westerwick, S., & Glynn, C. J. (2013). The Matilda Effect—Role Congruity Effects on Scholarly Communication: A Citation Analysis of Communication Research and Journal of Communication Articles. *Communication Research*, 40(1), 3–26. <https://doi.org/10.1177/0093650211418339>
- Knobloch-Westerwick, S., Glynn, C. J., & Huge, M. (2013). The Matilda Effect in Science Communication: An Experiment on Gender Bias in Publication Quality Perceptions and Collaboration Interest. *Science Communication*, 35(5), 603–625. <https://doi.org/10.1177/1075547012472684>

- Lakens, D. (2017). Equivalence Tests: A Practical Primer for t Tests, Correlations, and Meta-Analyses. *Social Psychological and Personality Science*, 8(4), 355–362.
<https://doi.org/10.1177/1948550617697177>
- Lakens, D., Scheel, A. M., & Isager, P. M. (2018). Equivalence Testing for Psychological Research: A Tutorial. *Advances in Methods and Practices in Psychological Science*, 1(2), 259–269. <https://doi.org/10.1177/2515245918770963>
- LeBel, E. P., Vanpaemel, W., Cheung, I., & Campbell, L. (2019). A Brief Guide to Evaluate Replications. *Meta-Psychology*, 3. <https://doi.org/10.15626/MP.2018.843>
- Ley, T. J., & Hamilton, B. H. (2008). The Gender Gap in NIH Grant Applications. *Science*, 322(5907), 1472–1474. <https://doi.org/10.1126/science.1165878>
- Lincoln, A. E., Pincus, S., Koster, J. B., & Leboy, P. S. (2012). The Matilda Effect in science: Awards and prizes in the US, 1990s and 2000s. *Social Studies of Science*, 42(2), 307–320. <https://doi.org/10.1177/0306312711435830>
- Lloyd, M. E. (1990). Gender Factors in Reviewer Recommendations for Manuscript Publication. *Journal of Applied Behavior Analysis*, 23(4), 539–543.
<https://doi.org/10.1901/jaba.1990.23-539>
- Madera, J. M., Hebl, M. R., & Martin, R. C. (2009). Gender and letters of recommendation for academia: Agentic and communal differences. *Journal of Applied Psychology*, 94(6), 1591–1599. <https://doi.org/10.1037/a0016539>
- Marsh, H. W., Bornmann, L., Mutz, R., Daniel, H.-D., & O'Mara, A. (2009). Gender Effects in the Peer Reviews of Grant Proposals: A Comprehensive Meta-Analysis Comparing Traditional and Multilevel Approaches. *Review of Educational Research*, 79(3), 1290–1326. <https://doi.org/10.3102/0034654309334143>

- Miller, D. I., & Wai, J. (2015). The bachelor's to Ph.D. STEM pipeline no longer leaks more women than men: a 30-year analysis. *Frontiers in Psychology*, 6.
<https://doi.org/10.3389/fpsyg.2015.00037>
- Milkman, K. L., Akinola, M., & Chugh, D. (2012). Temporal Distance and Discrimination: An Audit Study in Academia. *Psychological Science*, 23(7), 710–717.
<https://doi.org/10.1177/0956797611434539>
- Monroe, K., Ozyurt, S., Wrigley, T., & Alexander, A. (2008). Gender Equality in Academia: Bad News from the Trenches, and Some Possible Solutions. *Perspectives on Politics*, 6(2), 215–233. <https://doi.org/10.1017/S1537592708080572>
- Moss-Racusin, C. A., Dovidio, J. F., Brescoll, V. L., Graham, M. J., & Handelsman, J. (2012). Science faculty's subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41), 16474–16479.
- National Science Board. (2018). *Science & Engineering Indicators 2018*. Retrieved from <https://www.nsf.gov/statistics/2018/nsb20181/report/sections/academic-research-and-development/doctoral-scientists-and-engineers-in-academia>
- Nature Human Behaviour Editorial Board. (2018). Learning from replication. *Nature Human Behaviour*, 2(9), 601–601. <https://doi.org/10.1038/s41562-018-0441-1>
- Nelson, L. D., Simmons, J., & Simonsohn, U. (2018). Psychology's Renaissance. *Annual Review of Psychology*, 69(1), 511–534. <https://doi.org/10.1146/annurev-psych-122216-011836>
- Nittrouer, C. L., Hebl, M. R., Ashburn-Nardo, L., Trump-Steele, R. C. E., Lane, D. M., & Valian, V. (2018). Gender disparities in colloquium speakers at top universities. *Proceedings of the National Academy of Sciences*, 115(1), 104–108.
<https://doi.org/10.1073/pnas.1708414115>

- Open Science Collaboration. (2015). Estimating the reproducibility of psychological science. *Science*, 349(6251). <https://doi.org/10.1126/science.aac4716>
- Paulhus, D. L., Harms, P. D., Bruce, M. N., & Lysy, D. C. (2003). The over-claiming technique: Measuring self-enhancement independent of ability. *Journal of Personality and Social Psychology*, 84(4), 890–904. <https://doi.org/10.1037/0022-3514.84.4.890>
- Pohlhaus, J. R., Jiang, H., Wagner, R. M., Schaffer, W. T., & Pinn, V. W. (2011). Sex Differences in Application, Success, and Funding Rates for NIH Extramural Programs. *Academic Medicine*, 86(6), 759. <https://doi.org/10.1097/ACM.0b013e31821836ff>
- Primack, R. B., Ellwood, E., Miller-Rushing, A. J., Marrs, R., & Mulligan, A. (2009). Do gender, nationality, or academic age affect review decisions? An analysis of submissions to the journal Biological Conservation. *Biological Conservation*, 142(11), 2415–2418. <https://doi.org/10.1016/j.biocon.2009.06.021>
- Reuben, E., Sapienza, P., & Zingales, L. (2014). How stereotypes impair women's careers in science. *Proceedings of the National Academy of Sciences*, 111(12), 4403–4408.
- Richard, F. D., Bond, C. F. J., & Stokes-Zoota, J. J. (2003). One Hundred Years of Social Psychology Quantitatively Described. *Review of General Psychology*, 7(4), 331–363.
- Roberts, S. G., & Verhoef, T. (2016). Double-blind reviewing at EvoLang 11 reveals gender bias. *Journal of Language Evolution*, 1(2), 163–167. <https://doi.org/10.1093/jole/lzw009>
- Robnett, R. D. (2016). Gender Bias in STEM Fields: Variation in Prevalence and Links to STEM Self-Concept. *Psychology of Women Quarterly*, 40(1), 65–79. <https://doi.org/10.1177/0361684315596162>
- Rossiter, M. W. (1993). The Matthew Matilda Effect in Science. *Social Studies of Science*, 23(2), 325–341. <https://doi.org/10.1177/030631293023002004>

- 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(11), 1359–1366. <https://doi.org/10.1177/0956797611417632>
- Simons, D. J. (2014). The Value of Direct Replication. *Perspectives on Psychological Science*, 9(1), 76–80. <https://doi.org/10.1177/1745691613514755>
- Simonsohn, U. (2015). Small Telescopes: Detectability and the Evaluation of Replication Results. *Psychological Science*, 26(5), 559–569. <https://doi.org/10.1177/0956797614567341>
- Steele, J., James, J. B., & Barnett, R. C. (2002). Learning in a man's world: Examining the perceptions of undergraduate women in male-dominated academic areas. *Psychology of Women Quarterly*, 26(1), 46–50. <https://doi.org/10.1111/1471-6402.00042>
- Steinpreis, R. E., Anders, K. A., & Ritzke, D. (1999). The Impact of Gender on the Review of the Curricula Vitae of Job Applicants and Tenure Candidates: A National Empirical Study. *Sex Roles*, 41(7), 509–528. <https://doi.org/10.1023/A:1018839203698>
- Swim, J. K., Aikin, K. J., Hall, W. S., & Hunter, B. A. (1995). Sexism and racism: Old-fashioned and modern prejudices. *Journal of Personality and Social Psychology*, 68(2), 199.
- Teele, D. L., & Thelen, K. (2017). Gender in the Journals: Publication Patterns in Political Science. *PS: Political Science & Politics*, 50(2), 433–447. <https://doi.org/10.1017/S1049096516002985>
- The Carnegie Classification of Institutions of Higher Education. (n.d.). The Carnegie Classification of Institutions. Retrieved June 3, 2018, from <http://carnegieclassifications.iu.edu/>
- Tregenza, T. (2002). Gender bias in the refereeing process? *Trends in Ecology & Evolution*, 17(8), 349–350. [https://doi.org/10.1016/S0169-5347\(02\)02545-4](https://doi.org/10.1016/S0169-5347(02)02545-4)

- Trix, F., & Psenka, C. (2003). Exploring the Color of Glass: Letters of Recommendation for Female and Male Medical Faculty. *Discourse & Society*, 14(2), 191–220.
<https://doi.org/10.1177/0957926503014002277>
- Valian, V. (2004). Beyond Gender Schemas: Improving the Advancement of Women in Academia. *NWSA Journal*, 16(1), 207–220.
- van der Lee, R., & Ellemers, N. (2015). Gender contributes to personal research funding success in The Netherlands. *Proceedings of the National Academy of Sciences*, b(40), 12349–12353. <https://doi.org/10.1073/pnas.1510159112>
- Veldkamp, C. L. S., Hartgerink, C. H. J., Assen, M. A. L. M. van, & Wicherts, J. M. (2017). Who Believes in the Storybook Image of the Scientist? *Accountability in Research*, 24(3), 127–151. <https://doi.org/10.1080/08989621.2016.1268922>
- Wai, J., Lubinski, D., Benbow, C. P., & Steiger, J. H. (2010). Accomplishment in science, technology, engineering, and mathematics (STEM) and its relation to STEM educational dose: A 25-year longitudinal study. *Journal of Educational Psychology*, 102(4), 860–871.
<https://doi.org/10.1037/a0019454>
- Webb, T. J., O'Hara, B., & Freckleton, R. P. (2008). Does double-blind review benefit female authors? *Trends in Ecology & Evolution*, 23(7), 351–353.
<https://doi.org/10.1016/j.tree.2008.03.003>
- Wenneras, C., & Wold, A. (1997). Nepotism and sexism in peer-review. *Nature*, 387(6631), 341–343. <https://doi.org/10.1038/387341a0>
- West, J. D., Jacquet, J., King, M. M., Correll, S. J., & Bergstrom, C. T. (2013). The Role of Gender in Scholarly Authorship. *PLOS ONE*, 8(7), e66212.
<https://doi.org/10.1371/journal.pone.0066212>

- Wicherts, J. M., Veldkamp, C. L. S., Augusteijn, H. E. M., Bakker, M., van Aert, R. C. M., & van Assen, M. A. L. M. (2016). Degrees of Freedom in Planning, Running, Analyzing, and Reporting Psychological Studies: A Checklist to Avoid p-Hacking. *Frontiers in Psychology*, 7. <https://doi.org/10.3389/fpsyg.2016.01832>
- Williams, W. M., & Ceci, S. J. (2015). National hiring experiments reveal 2: 1 faculty preference for women on STEM tenure track. *Proceedings of the National Academy of Sciences*, 112(17), 5360–5365.
- Williams, J. C., Phillips, K. W., & Hall, E. V. (2016). Tools for change: Boosting the retention of women in the stem pipeline. *Journal of Research in Gender Studies*, 6(1), 11–75. <https://doi.org/10.22381/JRGS6120161>
- Zwaan, R. A., Pecher, D., Paolacci, G., Bouwmeester, S., Verkoeijen, P., Dijkstra, K., & Zeelenberg, R. (2018). Participant Nonnaiveté and the reproducibility of cognitive psychology. *Psychonomic Bulletin & Review*, 25(5), 1968–1972. <https://doi.org/10.3758/s13423-017-1348-y>