

When Journal Editors Play Favorites*

Remco Heesen[†]

December 5, 2016

Abstract

Should editors of scientific journals practice triple-anonymous reviewing? I consider two arguments in favor. The first says that insofar as editors' decisions are affected by information they would not have had under triple-anonymous review, an injustice is committed against certain authors. I show that even well meaning editors would commit this wrong and I endorse this argument.

The second argument says that insofar as editors' decisions are affected by information they would not have had under triple-anonymous review, it will negatively affect the quality of published papers. I distinguish between two kinds of biases that an editor might have. I show that one of them has a positive effect on quality and the other a negative one, and that the combined effect could be either positive or negative. Thus I do not endorse the second argument in general.

*Thanks to Kevin Zollman, Michael Strevens, Stephan Hartmann, Teddy Seidenfeld, Cailin O'Connor, Liam Bright, Shahar Avin, Jan-Willem Romeijn, an anonymous reviewer, and audiences at meetings of the Philosophy of Science Association in Atlanta and the Société de Philosophie des Sciences in Lausanne for valuable comments and discussion. This work was partially supported by the National Science Foundation under grant SES 1254291 and by an Early Career Fellowship from the Leverhulme Trust and the Isaac Newton Trust.

[†]Faculty of Philosophy, University of Cambridge, Sidgwick Avenue, Cambridge CB3 9DA, UK. Email: rdh51@cam.ac.uk.

However, I do endorse this argument for certain fields, for which I argue that the positive effect does not apply.

1 Introduction

Journal editors occupy an important position in the scientific landscape. By making the final decision on which papers get published in their journal and which papers do not, they have a significant influence on what work is given attention and what work is ignored in their field (Crane 1967).

In this paper I investigate the following question: should the editor be informed about the identity of the author when she is deciding whether to publish a particular paper? Under a single- or double-anonymous reviewing procedure, the editor has access to information about the author, whereas under a triple-anonymous reviewing procedure she does not.¹ So the question is: should journals practice triple-anonymous reviewing?

Two kinds of arguments have been given in favor of triple-anonymous reviewing. One focuses on the treatment of the author by the editor. On this kind of argument, revealing identity information to the editor will lead the editor to (partially) base her judgment on irrelevant information. This is unfair to the author, and is thus bad.

The second kind of argument focuses on the effect on the journal and its readers. Again, the idea is that the editor will base her judgment on identity information if given the chance to do so. But now the further claim is that as a result the journal will accept worse papers. After all, if a decision to accept or reject a paper is influenced by the editor's biases, this suggests that a departure has been made from a putative "objectively correct" decision. This harms the readers of the journal, and is thus bad.²

¹The relevant procedures are often called single-, double-, and triple-blind reviewing. I avoid this terminology as it has been criticized for being ableist (Tremain 2017, introduction).

²Hence, I distinguish between the effects of triple-anonymous reviewing on the author

This paper assesses these arguments. I distinguish between two different ways the editor’s judgment may be affected if the author’s identity is revealed to her. First, the editor may treat authors she knows differently from authors she does not know, a phenomenon I will call connection bias. Second, the editor may treat authors differently based on some aspect of their identity (e.g., their gender), which I will call identity bias. I make the following three claims.

My first claim is that connection bias actually benefits rather than harms the readers of the journal. This benefit is the result of a reduction in editorial uncertainty about the quality of submitted papers. I construct a model to show in a formally precise way how such a benefit might arise—surprisingly, no assumption that the scientists the editor knows are somehow “better scientists” is required—and I cite empirical evidence that such a benefit indeed does arise. However, this benefit only applies in certain fields; I argue that mathematics and parts of the humanities are excluded (section 2).

My second claim is that whenever connection bias or identity bias affects editorial decision, this constitutes an epistemic injustice in the sense of Fricker (2007) against the disadvantaged author. If the editor is to be (epistemically) just, she should prevent these biases from operating, which can be done through triple-anonymous reviewing. So I endorse an argument of the first of the two kinds I identified above: triple-anonymous reviewing is preferable because not doing so is unfair to authors (sections 3 and 4).

My third claim is that whether editorial biases harm the journal and its readers depends on a number of factors. Connection bias may benefit readers, whereas identity bias may harm them. Whether there is an overall benefit or harm depends on the strength of the editor’s identity bias, the relative sizes of the different groups, and other factors, as I illustrate using the model. As

and on the readers of the journal. This reflects a growing understanding that in order to study the social epistemology of science, what is good for an individual inquirer must be distinguished from what is good for the wider scientific community (Kitcher 1993, Strevens 2003, Mayo-Wilson et al. 2011).

a result I do not in general endorse the second kind of argument, that triple-anonymous reviewing is preferable because readers of the journal are harmed otherwise. However, I do endorse this argument for fields like mathematics, where I claim that the benefits of connection bias do not apply (section 5).

Zollman (2009) has studied the effects of different editorial policies on the number of papers published and the selection criteria for publication, but he does not focus specifically on the editor's decisions. Economists have studied models in which editorial decisions play an important role (Ellison 2002, Faria 2005, Besancenot et al. 2012), but they have not been concerned with biases the editor may be subject to. Other economists have done empirical work investigating the differences between papers with and without an author-editor connection (Laband and Piette 1994, Medoff 2003, Smith and Dombrowski 1998, more on this later), but they do not provide a model that can explain these differences. This paper thus fills a gap in the literature.

In particular, I compare double- and triple-anonymous reviewing as opposed to single- and double-anonymous reviewing. The latter comparison has been studied extensively, see, e.g., Blank (1991) for a prominent empirical study and Snodgrass (2006) and Lee et al. (2013, especially pp. 10–11) for literature reviews. In contrast, I know of almost no empirical or theoretical work directly comparing double- and triple-anonymous reviewing (one exception is Lee and Schunn 2010, p. 7).

While I focus on comparing double- and triple-anonymous review, some of what I say may carry over to the context of comparing single- and double-anonymous review. In section 5 I comment briefly on the extent to which the formal model I present applies in the context of comparing single- and double-anonymous review. However, I leave it to the reader to judge to what extent the arguments I make on the basis of the model carry over.

2 A Model of Connection Bias

As mentioned, journal editors have a certain measure of power in a scientific community because they decide which papers get published.³ An editor could use this power to the benefit of her friends or colleagues, or to promote certain subfields or methodologies over others. This phenomenon has been called *editorial favoritism*.

Much of the evidence for favoritism is anecdotal, but some systematic studies exist as well. Bailey et al. (2008a,b) find that academics believe editorial favoritism to be fairly prevalent, with a nonnegligible percentage claiming to have perceived it firsthand. Hull (1988, chapter 9) finds a limited degree of favoritism towards authors sharing the editor’s views in his detailed study of reviewing practices at the journal *Systematic Zoology*. And Laband (1985) and Piette and Ross (1992) find that, controlling for citation impact and various other factors, papers whose author has a connection to the journal editor are allocated more journal pages than papers by authors without such a connection.⁴

In this paper, I refer to the phenomenon that editors are more likely to accept papers from authors they know than papers from authors they do not know as *connection bias*.

Academics tend to disapprove of this behavior (Sherrell et al. 1989, Bailey et al. 2008a,b). In both studies by Bailey et al., in which subjects were asked to rate the seriousness of various potentially problematic behaviors by editors and reviewers, this disapproval was shown (using a factor analysis) to be part of a general and strong disapproval of “selfish or cliquish acts” in the

³Different journals may have different policies, such as one in which associate editors make the final decision for papers in their (sub)field. Here, I simply define “the editor” to be whomever makes the final decision whether to publish a particular paper.

⁴Here, page allocation is used as a proxy for journal editors’ willingness to push the paper. The more obvious variable to use here would be whether or not the paper is accepted for publication. Unfortunately, there are no empirical studies which measure the influence of author-editor relationships on acceptance decisions directly. Presumably this is because information about rejected papers is usually not available in these studies.

peer review process.⁵ Thus it appears that the reason for the disapproval of connection bias is that it shows the editor acting on private interests, rather than displaying the disinterestedness that is the norm in science (Merton 1942).

On the other hand, there is some evidence that connection bias improves the overall quality of accepted papers (Laband and Piette 1994, Medoff 2003, Smith and Dombrowski 1998). Does this mean scientists are misguided in their disapproval?

In this section, I use a formal model to show that editors may display connection bias even if their only goal is to accept the best papers, and that this may improve quality, consistent with Laband and Piette's, Medoff's, and Smith and Dombrowski's findings. Note that in this section I discuss connection bias only. Subsequent sections discuss identity bias.

Consider a simplified scientific community. Each scientist produces a paper and submits it to the community's only journal which has one editor.

Some papers are more suitable for publication than others. I assume that this suitability can be measured on a single numerical scale. For convenience I call this the *quality* of the paper. However, I remain neutral on how this notion should be interpreted, e.g., as an objective measure of the epistemic value of the paper, or as the number of times the paper would be cited in future papers if it was published, or as the average subjective value each member of the scientific community would assign to it if they read it (see Bright 2015, for more on potential difficulties with interpreting the notion of quality).

⁵This evidence conflicts to some extent with other survey findings. If connection bias was a serious worry for working scientists, one would expect them to rank knowing the editor and the composition of the editorial board more generally among the most important factors in deciding where to submit their papers. But Ziobrowski and Gibler (2000) find that this is not the case (these factors are ranked twelfth and sixteenth in a list of sixteen potentially relevant factors in their survey). In a similar survey by Mackie (1998, chapter 4), twenty percent of authors indicated that knowing the editor and/or her preferences is an important consideration in deciding where to submit a paper.

Crucially, the editor does not know the quality of the paper at the time it is submitted. The aim of this section is to show how uncertainty about quality can lead to connection bias. To make this point as starkly as possible, I assume that the editor cares only about quality, i.e., she makes an estimate of the quality of a paper and publishes those and only those papers whose quality estimate is high.

Let q_i be the quality of the paper submitted by scientist i . q_i is modeled as a random variable to reflect uncertainty about quality. Since some scientists are more likely to produce high quality papers than others, the mean μ_i of this random variable may be different for each scientist. I assume that quality follows a normal distribution with fixed variance, writing $q_i | \mu_i \sim N(\mu_i, \sigma_{in}^2)$ for “ q_i given μ_i follows a normal distribution with mean μ_i and variance σ_{in}^2 ”. The subscript in indicates that this is the variance in individual papers by the same author.

The assumptions of normality and fixed variance are made primarily to keep the mathematics simple. Below I make similar assumptions on the distribution of average quality in the scientific community and the distribution of reviewers’ estimates of the quality of a paper. I see no reason to expect the results below to be different when these assumptions are changed.

If the editor knows scientist i , she has some prior information on the average quality of scientist i ’s work. This is reflected in the model by assuming that the editor knows the value of μ_i . In contrast, the editor is uncertain about the average quality of the work of scientists she does not know. All she knows is the distribution of average quality in the larger scientific community, which I also assume to be normal: $\mu_i \sim N(\mu, \sigma_{sc}^2)$.

Note that I assume the scientific community to be homogeneous: average paper quality follows the same distribution in the two groups of scientists (those known to the editor and those not known to the editor). If I assumed instead that scientists known to the editor write better papers on average the results would be qualitatively similar to those I present below. If scientists

known to the editor write worse papers on average this would affect my results. However, since most journal editors are relatively central figures in their field (Crane 1967), this seems implausible for most cases.

The editor’s prior for the quality of a paper submitted by some scientist i reflects this difference in information. If she knows the scientist she knows the value of μ_i , and so her prior is $\pi(q_i | \mu_i) \sim N(\mu_i, \sigma_{in}^2)$. If the editor does not know scientist i she is uncertain about μ_i . Integrating out this uncertainty yields a prior $\pi(q_i) \sim N(\mu, \sigma_{in}^2 + \sigma_{sc}^2)$ for the quality of scientist i ’s paper.

When the editor receives a paper she sends it out for review. The reviewer provides an estimate r_i of the paper’s quality which is again a random variable. I assume that the reviewer’s report is unbiased, i.e., its mean is the actual quality q_i of the paper. Once again I use a normal distribution to reflect uncertainty: $r_i | q_i \sim N(q_i, \sigma_{rv}^2)$.⁶

The editor uses the information from the reviewer’s report to update her beliefs. I assume that she does this by conditioning on r_i . Thus, her posterior for the quality of scientist i ’s paper is $\pi(q_i | r_i)$ if she does not know the author, and $\pi(q_i | r_i, \mu_i)$ if she does.

The posterior distributions are themselves normal distributions whose mean is a weighted average of r_i and the prior mean (see proposition 5 in appendix A). I write μ_i^U for the mean of the posterior distribution if the editor does not know scientist i and μ_i^K for the posterior mean if she does.

When does the editor choose to publish a paper? Here I assume that

⁶The reviewer’s report could reflect the opinion of a single reviewer, or the averaged opinion of multiple reviewers. The editor could even act as a reviewer herself, in which case the report reflects her findings which she has to incorporate in her overall beliefs about the quality of the paper. The assumption I make in the text can be used to cover any of these scenarios, as long as a given journal is fairly consistent in the number of reviewers used. Some journals may use different numbers of reviewers for different papers (potentially affecting the variance if more reviewers give more accurate information than fewer) or employ reviewers in different roles (e.g., one reviewer to assess technical aspects of the paper and one reviewer to assess non-technical aspects). My model does not apply to journals where these differences correlate with the existence or absence of a connection between editor and author.

she publishes any paper whose (posterior) expected quality is above some threshold q^* . So a paper written by a scientist unknown to the editor is published if $\mu_i^U > q^*$ and a paper written by a scientist known to the editor is published if $\mu_i^K > q^*$. Other standards could be used: risk-averse standards might require high (greater than 50%) confidence that the paper is above some threshold. In this case the threshold value needs to be lowered to keep the journal's acceptance rate constant. For my purposes here this makes little difference.

The first result of this section establishes the existence of connection bias in the model (see appendix A for proofs of all theorems). It says that the editor is more likely to publish a paper written by an arbitrary author she knows than a paper written by an arbitrary author she does not know, whenever $q^* > \mu$ (for any positive value of σ_{sc}^2 and σ_{rv}^2). Since $q^* = \mu$ would mean that exactly half of all papers gets published, the condition amounts to a requirement that the journal's acceptance rate is less than 50%. This is true of most reputable journals in most fields (physics being a notable exception).

Theorem 1 (Connection Bias). *If $q^* > \mu$, $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$, the acceptance probability for authors known to the editor is higher than the acceptance probability for authors unknown to the editor, i.e., $\Pr(\mu_i^K > q^*) > \Pr(\mu_i^U > q^*)$.*

Theorem 1 shows that in my model any journal with an acceptance rate lower than 50% will be seen to display connection bias. Thus I have established the surprising result that an editor who cares only about the quality of the papers she publishes may end up publishing more papers by her friends and colleagues than by scientists unknown to her, even if her friends and colleagues are not, as a group, better scientists than average.⁷

⁷The model presented in this section is formally similar to that of Miller (1994) and Borsboom et al. (2008). They assume genuine differences in average quality between groups, so a result like theorem 1 is true but unsurprising in their models.

Why does this surprising result hold? The distribution of the posterior mean μ_i^U has lower variance than the distribution of μ_i^K (see proposition 6 in appendix A). That is, the variance of μ_i^U is lower in an “objective” sense; this is not a claim about the editor’s subjective uncertainty about her judgment. This is because μ_i^U is a weighted average of μ and r_i , keeping it relatively close to the overall mean μ compared to μ_i^K , which is a weighted average of μ_i and r_i (which tend to differ from μ in the same direction).

Note that the result assumes that scientists known to the editor and scientists unknown to the editor are held to the same “standard” (the threshold q^*). Alternatively, the editor might enforce equal acceptance rates for the two groups. This would be formally equivalent to raising the threshold for known scientists (and/or lowering the threshold for unknown scientists).

What I have shown so far is that an editor who uses information about the average quality of papers produced by scientists she knows in her acceptance decisions will find that scientists she knows produce on average more papers that meet her quality threshold. This is a subjective statement: the editor believes that more papers by scientists she knows meet her threshold. Does this translate into an objective effect?

In order to answer this question I need to compare the average quality of accepted papers. More formally, I want to compare the expected value of the quality of a paper, conditional on meeting the publication threshold, given that the author is either known to the editor or not.

Theorem 2 (Positive Effect of Connection Bias). *If $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$, the average quality of accepted papers from authors known to the editor is higher than the average quality of accepted papers from authors unknown to the editor, i.e., $\mathbb{E}[q_i \mid \mu_i^K > q^*] > \mathbb{E}[q_i \mid \mu_i^U > q^*]$.*

The editor’s knowledge of the average quality of papers written by scientists she knows makes it such that among those scientists relatively many whose papers are accepted have relatively high average quality. Since this

correlates with paper quality the average quality of accepted papers in this group is relatively high, yielding theorem 2.

The theorem shows that the editor can use the extra information she has about scientists she knows to improve the average quality of the papers published in her journal. The surprising result, then, is that the editor's connection bias actually benefits rather than harms the readers of the journal. In other words, the editor can use her connections to "identify and capture high-quality papers", as Laband and Piette (1994) suggest.

To what extent does this show that the connection bias observed in reality is the result of editors capturing high-quality papers, as opposed to editors using their position of power to help their friends? At this point the model is seen to yield an empirical prediction. If connection bias is (primarily) due to capturing high-quality papers, the quality of papers by authors the editor knows should be higher than average, as shown in the model. If, on the other hand, connection bias is (primarily) a result of the editor accepting for publication papers written by authors she knows even though they do not meet the quality standards of the journal, then the quality of papers by authors the editor knows should presumably be lower than average.

If subsequent citations are a good indication of the quality of a paper,⁸ a simple regression can test whether accepted papers written by authors with an author-editor connection have higher or lower average quality than papers without such a connection. This empirical test has been carried out a number of times, and the results favor the hypothesis that editors use their connections to improve the quality of published papers (Laband and Piette 1994, Smith and Dombrowski 1998, Medoff 2003).⁹

⁸Recall that I have remained neutral on how the notion of quality should be interpreted. If quality is simply defined as "the number of citations this paper would get if it were published" the connection between quality and citations is obvious. Even on other interpretations of quality, citations have frequently been viewed as a good proxy measure (Cole and Cole 1967, 1968, Medoff 2003). This practice has been defended by Cole and Cole (1971) and Clark (1957, chapter 3), and criticized by Lindsey (1989) and Heesen (forthcoming).

⁹Laband and Piette and Medoff focus on economics journals and Smith and Dom-

Note that in the above (qualitative) results, nothing depends on the sizes of the variances σ_{in}^2 , σ_{sc}^2 , and σ_{rv}^2 . The values of the variances do matter when the acceptance rate and average quality of papers are compared quantitatively. For example, reducing σ_{rv}^2 (making the reviewer's report more accurate) reduces the differences in the acceptance rate and average quality of papers.

Note also that the results depend on the assumption that σ_{sc}^2 and σ_{rv}^2 are positive. What is the significance of these assumptions?

If $\sigma_{rv}^2 = 0$, i.e., if there is no variance in the reviewer's report, the reviewer's report describes the quality of the paper with perfect accuracy. In this case the "extra information" the editor has about authors she knows is not needed, and so there is no difference in acceptance rate or average quality based on whether the editor knows the author. But it seems unrealistic to expect reviewer's reports to be this accurate.

If $\sigma_{sc}^2 = 0$ there is either no difference in the average quality of papers produced by different authors, or learning the identity of the author does not tell the editor anything about the expected quality of that scientist's work. In this case there is no value to the editor (with regard to determining the quality of the submitted paper) in learning the identity of the author. So here there is also no difference in acceptance rate or average quality based on whether the editor knows the author.

Under what circumstances should the identity of the author be expected to tell the editor something useful about the quality of a submitted paper? This seems to be most obviously the case in the lab sciences. The identity of the author, and hence the lab at which the experiments were performed, can increase or decrease the editor's confidence that the experiments were performed correctly, including all the little checks and details that are impossible to describe in a paper. In such cases, "the reader must rely on the author's

browski on accounting journals. Further research would be valuable to see whether these results generalize, especially to the natural sciences and the humanities.

(and perhaps referee’s) testimony that the author really performed the experiment exactly as claimed, and that it worked out as reported” (Easwaran 2009, p. 359).

But in other fields, in particular mathematics and those parts of the humanities that focus on abstract arguments, there is no need to rely on the author’s reputation. This is because in these fields the paper itself is the contribution, so it is possible to judge papers in isolation of how or by whom they were created (Easwaran 2009). And in fact there exists a norm that this is how they should be judged: “Papers will rely only on premises that the competent reader can be assumed to antecedently believe, and only make inferences that the competent reader would be expected to accept on her own consideration.” (Easwaran 2009, p. 354).

Arguably then, the epistemic advantage conferred by revealing identity information about the author to the editor applies only in certain fields. The relevant fields are those where part of the information in the paper is conferred on the authority of testimony. In mathematics and parts of the humanities, where a careful reading of a paper itself constitutes a reproduction of its argument, there is no relevant information to be learned from the identity of the author (i.e., $\sigma_{sc}^2 = 0$). Or at least the publishing norms in these fields suggest that their members believe this to be the case.

3 Connection Bias As an Epistemic Injustice

The previous section discussed a formal model of editorial uncertainty about paper quality. I first established the existence of connection bias in this model. Then I showed that connection bias benefits the readers of the journal, insofar as readers care about the quality of accepted papers. Despite this benefit to readers, I claim that connection bias is unfair to authors. In this section I argue this claim by appealing to the concept of *epistemic injustice*, as developed by Fricker (2007).

The type of epistemic justice that is relevant here is *testimonial injustice*. Fricker (2007, pp. 17–23) defines a testimonial injustice as a case where a speaker suffers a credibility deficit for which the hearer is ethically and epistemically culpable, rather than being due to innocent error.

Testimonial injustices may arise in various ways. Fricker is particularly interested in what she calls “the central case of testimonial injustice” (Fricker 2007, p. 28). This kind of injustice results from a *negative identity-prejudicial stereotype*, which is defined as follows:

A widely held disparaging association between a social group and one or more attributes, where this association embodies a generalization that displays some (typically, epistemically culpable) resistance to counter-evidence owing to an ethically bad affective investment. (Fricker 2007, p. 35)

Because the stereotype is widely held, it produces *systematic* testimonial injustice: the relevant social group will suffer a credibility deficit in many different social spheres.

Applying this to the phenomenon of connection bias, it is clear that this is not an instance of the central case of testimonial injustice. This would require some negative stereotype associated with scientists unknown to the editor (as a group) which does not normally exist. So I set the central case aside for now and focus on the question whether connection bias can produce (non-central cases of) testimonial injustice.

How are individual scientists affected by the differential acceptance rates established in section 2? For scientist i , the probability of acceptance given the average quality of her papers μ_i denotes the long-run average proportion of her papers that will be accepted (assuming she submits all her papers to the journal).

Theorem 3 (Acceptance Rate for Individual Authors). *Assume $\sigma_{sc}^2 > 0$ and $\sigma_{rv}^2 > 0$. The acceptance rate for author i (with average quality μ_i) is higher if the editor knows her if and only if μ_i exceeds a weighted average of μ and q^* :*

$$\Pr(\mu_i^K > q^* \mid \mu_i) \geq \Pr(\mu_i^U > q^* \mid \mu_i) \quad \text{iff} \quad \mu_i \geq \frac{\sigma_{in}^2}{\sigma_{in}^2 + \sigma_{sc}^2} \mu + \frac{\sigma_{sc}^2}{\sigma_{in}^2 + \sigma_{sc}^2} q^*.$$

Because theorem 3 states a biconditional, the strict version is true as well, i.e., if the editor knows scientist i she is strictly better off if and only if μ_i strictly exceeds the weighted average.

Note that regardless of the values of the variances, any scientist whose average quality exceeds the threshold value ($\mu_i \geq q^*$) benefits from connection bias. Conversely, a scientist of below average quality ($\mu_i \leq \mu$) is actually worse off if the editor knows her.¹⁰

Consider what this theorem says for a particular scientist i who is unknown to the editor and whose average quality μ_i strictly exceeds the weighted average. Her best papers are accepted regardless of her lack of connection to the editor, and her worst papers are rejected regardless. But there exists a middle range of papers that are actually rejected but that would have been accepted if only the editor knew her.

In Fricker's terminology, scientist i suffers from a credibility deficit: fewer of her papers are considered credible by the editor than would have been considered credible if the editor knew her.

Is this credibility deficit suffered by scientist i ethically and epistemically culpable on the part of the editor? On the one hand, the editor is simply making maximal use of the information available to her. It just so happens that she has more information about scientists she knows than about others. But that is hardly the editor's fault. Is it incumbent upon her to get to know the work of every scientist who submits a paper?

This may well be too much to ask. But an alternative option is to remove all information about the authors of submitted papers. This can be done by using a triple-anonymous reviewing procedure, in which the editor

¹⁰These claims assume that $q^* \geq \mu$. If $q^* < \mu$, the scientist benefits if $\mu_i \geq \mu$ and is harmed if $\mu_i \leq q^*$.

does not know the identity of the author, and hence is prevented from using information about scientists she knows in her evaluation.

I conclude that the editor is ethically and epistemically culpable for credibility deficits suffered by scientists unknown to the editor whose average quality exceeds the weighted average specified in theorem 3, and hence testimonial injustices are committed against such authors when a double-anonymous reviewing procedure is used. A similar epistemic injustice occurs for scientists known to the editor whose average quality is below the weighted average, as such authors would prefer that the editor not use information she has about their average quality.

It is worth noting explicitly which scientists are better or worse off in terms of acceptance rates if a triple-anonymous procedure is introduced. If the acceptance threshold q^* is held constant¹¹, nothing changes for scientists unknown to the editor. Scientists known to the editor will see their acceptance rate go down if their average quality exceeds the weighted average specified in theorem 3, and up otherwise. The overall acceptance rate of the journal will go down.

So the group that I based my argument on (unknown scientists of high average quality) is not necessarily made better off by switching to triple-anonymous reviewing. The argument for triple-anonymous reviewing given in this section is not about benefitting one group of scientists or harming another: rather, it is about fairness. Under a triple-anonymous procedure, at least all scientists are treated equally: any scientist who writes a paper of a given quality has the same chance of seeing that paper accepted. Whereas under a double-anonymous procedure, scientists are treated unfairly in that their acceptance rates may differ based only on an epistemically irrelevant

¹¹Things are slightly more subtle if the overall acceptance rate of the journal is held constant instead. The threshold will go down, and hence all scientists unknown to the editor, as well as known scientists of low average quality, will see their acceptance rates go up. Known scientists of high average quality must then necessarily see their acceptance rate go down as a group, but the individual effect is non-trivial to determine as they no longer benefit from connection bias but they do benefit from the decreased threshold.

characteristic (knowing the editor).

I conclude that while journal readers may benefit from connection bias, it involves unfair treatment of authors. Because this unfair treatment takes the form of an epistemic injustice, which involves both ethically and epistemically culpable behavior, connection bias has both an epistemic benefit (to readers) and a cost (to the author). It would be a misinterpretation of my analysis, then, to conclude that connection bias is epistemically good but ethically bad.

4 Identity Bias

So far, I have assumed that connection bias is the only bias journal editors display. The literature on implicit bias suggests that this is not true. For example, “[i]f submissions are not anonymous to the editor, then the evidence suggests that women’s work will probably be judged more negatively than men’s work of the same quality” (Saul 2013, p. 45). Evidence for this claim is given by Wennerås and Wold (1997), Valian (1999, chapter 11), Steinfeld et al. (1999), Budden et al. (2008), and Moss-Racusin et al. (2012).¹² So women scientists are at a disadvantage simply because of their gender identity. Similar biases exist based on other irrelevant aspects of scientists’ identity, such as race or sexual orientation (see Lee et al. 2013, for a critical survey of various biases in the peer review system). As Crandall (1982, p. 208) puts it: “The editorial process has tended to be run as an informal, old-boy network which has excluded minorities, women, younger researchers, and those from lower-prestige institutions”.

I use *identity bias* to refer to these kinds of biases. I now complicate the model of section 2 to include identity bias. I then argue that allowing

¹²These citations show that the work of women in academia is undervalued in various ways. None of them focus specifically on editor evaluations, but they support Saul’s claim unless it is assumed that journal editors as a group are significantly less biased than other academics.

the editor’s decisions to be influenced by identity bias is unfair to authors, analogous to the argument of the previous section.

I incorporate identity bias in the model by assuming the editor consistently undervalues members of one group (and overvalues the others). More precisely, she believes the average quality of papers produced by any scientist i from the group she is biased against to be lower than it really is by some constant quantity ε . Conversely, the average quality of papers written by any scientist not belonging to this group is raised by δ .¹³ So the editor has a different prior for the two groups; I use π_A to denote her prior for the quality of papers written by scientists she is biased against, and π_F for her prior for scientists she is biased in favor of.

As before, the editor may know a given scientist or not. So there are now four groups. If scientist i is known to the editor and belongs to the stigmatized group the editor’s prior distribution on the quality of scientist i ’s paper is $\pi_A(q_i | \mu_i) \sim N(\mu_i - \varepsilon, \sigma_{in}^2)$. If scientist i is known to the editor but is not in the stigmatized group the prior is $\pi_F(q_i | \mu_i) \sim N(\mu_i + \delta, \sigma_{in}^2)$. If scientist i is not known to the editor and is in the stigmatized group the prior is $\pi_A(q_i) \sim N(\mu - \varepsilon, \sigma_{in}^2 + \sigma_{sc}^2)$. And if scientist i is not known to the editor and not in the stigmatized group the prior is $\pi_F(q_i) \sim N(\mu + \delta, \sigma_{in}^2 + \sigma_{sc}^2)$.¹⁴

After the reviewer’s report comes in the editor updates her beliefs about the quality of the paper. This yields posterior distributions $\pi_A(q_i | r_i, \mu_i)$,

¹³This is a simplifying assumption: one could imagine having biases against multiple groups of different strengths, biases whose strength has some random variation, or biases which intersect in various ways (Collins and Chepp 2013, Bright et al. 2016). However, the assumption in the main text suffices for my purposes. It should be fairly straightforward to extend my results to more complicated cases like the ones just described.

¹⁴Note that I assume that the editor displays bias against scientists in the stigmatized group regardless of whether she knows them or not. Under a reviewing procedure that is not triple-anonymous, the editor learns at least the name and affiliation of any scientist who submits a paper. This information is usually sufficient to determine with reasonable certainty the scientist’s gender. So at least for gender bias it seems reasonable to expect the editor to display bias even against scientists she does not know. Conversely, because negative identity-prejudicial stereotypes can work unconsciously, it does not seem reasonable to expect that the editor can withhold her bias from scientists she knows.

$\pi_F(q_i | r_i, \mu_i)$, $\pi_A(q_i | r_i)$, and $\pi_F(q_i | r_i)$, with posterior means μ_i^{KA} , μ_i^{KF} , μ_i^{UA} , and μ_i^{UF} , respectively. As before, the paper is published if the posterior mean exceeds the threshold q^* . This yields the unsurprising result that the editor is less likely to publish papers by scientists she is biased against.

Theorem 4 (Identity Bias). *If $\varepsilon + \delta > 0$, $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$, the acceptance probability for authors the editor is biased against is lower than the acceptance probability for authors the editor is biased in favor of (keeping fixed whether or not the editor knows the author). That is,*

$$\Pr(\mu_i^{KA} > q^*) < \Pr(\mu_i^{KF} > q^*) \quad \text{and} \quad \Pr(\mu_i^{UA} > q^*) < \Pr(\mu_i^{UF} > q^*).$$

Theorem 4 establishes the existence of identity bias in the model: authors that the editor is biased against are less likely to see their paper accepted than other authors.

Any time a paper is rejected because of identity bias (i.e., the paper would have been accepted if the relevant part of the author’s identity had been different, all else being equal), a testimonial injustice occurs. Moreover, here the editor is culpable for having these biases.

Testimonial injustices resulting from identity bias can be instances of the central case of testimonial injustice, in which the credibility deficit results from a negative identity-prejudicial stereotype. The evidence suggests that negative identity-prejudicial stereotypes affect the way people (not just men) judge women’s work, even when the person judging does not consciously believe in these stereotypes. Moreover, those who think highly of their ability to judge work objectively and/or are primed with objectivity are affected more rather than less (Uhlmann and Cohen 2007, Stewart and Payne 2008, p. 1333). Similar claims plausibly hold for biases based on race or sexual orientation. Biases based on academic affiliation are not usually due to negative identity-prejudicial stereotypes, as these do not generally affect other aspects of the scientist’s life.

So both connection bias and identity bias are responsible for injustices against authors. This is one way to spell out the claim that it is unfair to authors when journal editors do not use a triple-anonymous reviewing procedure. This constitutes the first kind of argument for triple-anonymous reviewing which I mentioned in the introduction, and which I endorse based on these considerations.

5 The Tradeoff Between Connection Bias and Identity Bias

The second kind of argument I mentioned in the introduction claims that failing to use triple-anonymous reviewing harms the journal and its readers, because it would lower the average quality of accepted papers. In section 2 I argued that connection bias actually has the opposite effect: it increases average quality. The introduction of identity bias complicates the picture, as identity bias generally lowers the average quality of accepted papers. This raises the question whether the combined effect of connection bias and identity bias is positive or negative. In this section I show that there is no general answer to this question.

I compare the average quality of accepted papers under a procedure subject to connection bias and identity bias to that under a triple-anonymous reviewing procedure. Under this procedure, the editor's prior distribution for the quality of any submitted paper is $\pi(q_i) \sim N(\mu, \sigma_{in}^2 + \sigma_{sc}^2)$, i.e., the prior I used in section 2 when the author was unknown to the editor. Hence the posterior is $\pi(q_i | r_i)$ with mean μ_i^U , the probability of acceptance is $\Pr(\mu_i^U > q^*)$ and the average quality of accepted papers is $\mathbb{E}[q_i | \mu_i^U > q^*]$. As a result, the editor displays neither connection bias nor identity bias.

In contrast, the double-anonymous reviewing procedure is subject to connection bias and identity bias. The overall probability that a paper is accepted under this procedure depends on the relative sizes of the four groups.

I use p_{KA} to denote the fraction of scientists known to the editor that she is biased against, p_{KF} for the fraction known to the editor that she is biased in favor of, p_{UA} for unknown scientists biased against, and p_{UF} for unknown scientists biased in favor of. These fractions are nonnegative and sum to one.

Let A_i denote the event that scientist i 's paper is accepted under the double-anonymous procedure. The overall probability of acceptance is

$$\begin{aligned} \Pr(A_i) &= p_{KA} \Pr(\mu_i^{KA} > q^*) + p_{KF} \Pr(\mu_i^{KF} > q^*) \\ &\quad + p_{UA} \Pr(\mu_i^{UA} > q^*) + p_{UF} \Pr(\mu_i^{UF} > q^*). \end{aligned}$$

The average quality of accepted papers can then be written as $\mathbb{E}[q_i | A_i]$, which I compare to $\mathbb{E}[q_i | \mu_i^U > q^*]$.¹⁵

In the remainder of this section I assume that the editor's biases are such that she believes the average quality of all submitted papers to be equal to the overall average μ . In other words, her bias against the stigmatized group is canceled out on average by her bias in favor of those not in the stigmatized group, weighted by the relative sizes of those groups:

$$(p_{KA} + p_{UA}) \varepsilon = (p_{KF} + p_{UF}) \delta.$$

I use the above equation to fix the value of δ , reducing the number of free parameters by one. This is a kind of commensurability requirement for the two procedures because it guarantees that the editor perceives the average quality of submitted papers to be μ regardless of which reviewing procedure is used.

As far as I can tell there are no interesting general conditions on the parameter values that determine whether the double-anonymous procedure

¹⁵Expressions for $\Pr(A_i)$ and $\mathbb{E}[q_i | A_i]$ using only the parameter values and standard functions are given in proposition 13 in appendix A. These expressions are used to generate the numerical results below.

or the triple-anonymous procedure will lead to a higher average quality of accepted papers. The question I will explore now, using some numerical examples, is how biased the editor needs to be for the epistemic costs of her identity bias to outweigh the epistemic benefits resulting from connection bias.

In order to generate numerical data values have to be chosen for the parameters. First I set $\mu = 0$ and $q^* = 2$. Since quality is an interval scale in this model, these choices are arbitrary. For the variances σ_{in}^2 (of the quality of individual papers), σ_{sc}^2 (of the average quality of authors), and σ_{rv}^2 (of the accuracy of the reviewer’s report), I choose a “small” and a “large” value (1 and 4 respectively).

For the sizes of the four groups, I assume that there is no correlation between whether the editor knows an author and whether the editor has a bias against that author (so, e.g., the percentage of women among scientists the editor knows is equal to the percentage of women among scientists the editor does not know). I consider two cases for the editor’s identity bias: either she is biased against exactly half of all authors or the group she is biased against is a 30 % minority.¹⁶ Similarly, I consider the case in which the editor knows half of all scientists submitting papers, and the case in which the editor knows 30 % of them.

As a result, there are 32 possible settings of the parameters (2^3 choices for the variances times 2^2 choices for the group sizes).

It follows from theorem 2 that when $\varepsilon = 0$ the double-anonymous procedure helps rather than harms the readers of the journal by increasing average quality relative to the triple-anonymous procedure. If ε is positive but relatively small, this remains true, but when ε is relatively big, the double-

¹⁶Bruner and O’Connor (forthcoming) note that certain dynamics in academic life can lead to identity bias against groups as a result of the mere fact that they are a minority. Here I consider both the case where the stigmatized group is a minority (and is possibly stigmatized as a result of being a minority, as Bruner and O’Connor suggest) and the case where it is not (and so presumably the negative identity-prejudicial stereotype has some other source).

anonymous procedure harms the readers. This is because the average quality of published papers under the double-anonymous procedure decreases continuously as ε increases.

The interesting question, then, is where the turning point lies. How big does the editor’s bias need to be in order for the negative effects of identity bias on quality to cancel out the positive effects of connection bias?

I determine the value of ε for which the average quality of published papers under the double-anonymous procedure and the triple-anonymous procedure is the same for each of the 32 cases. Figure 1 reports these numbers. I plot them against the acceptance rate that the triple-anonymous procedure would have for those values of the parameters, i.e., $\Pr(\mu_i^U > q^*)$. The bias ε is measured in “quality points” (for reference: since $\mu = 0$ and $q^* = 2$, a paper needs to be two quality points above average to be accepted).

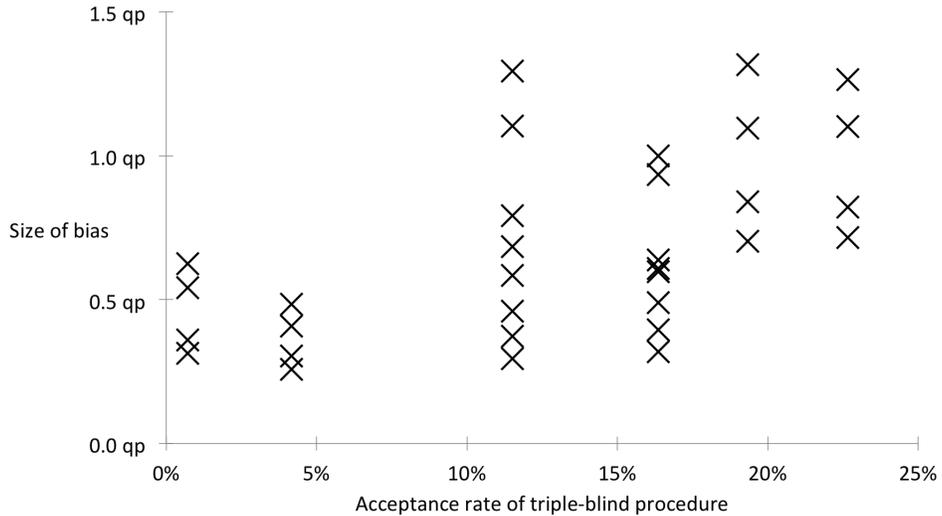


Figure 1: The minimum size of the editor’s bias such that the quality costs of the double-anonymous procedure outweigh its benefits (measured in “quality points”), in 32 cases, plotted as a function of the acceptance rate of the corresponding triple-anonymous procedure.

Since quality points do not have a clear interpretation outside the context of the model, I use the values of ε shown in figure 1 to calculate the average

rate of acceptance of papers from authors the editor is biased against and the average rate of acceptance of papers from authors the editor is biased in favor of.¹⁷ The difference between these numbers gives an indication of the size of the editor’s bias: it measures (in percentage points, abbreviated pp) how many more papers the editor accepts from authors she is biased in favor of, compared to those she is biased against.

This difference is reported for the 32 cases in figure 2. Even with this small sample of 32 cases, a large variation of results can be observed. I illustrate this by looking at two cases in detail.

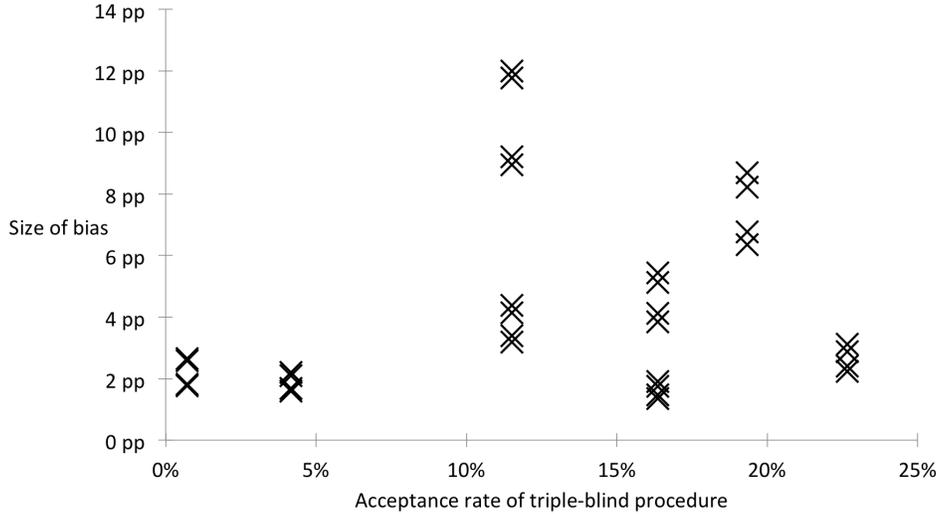


Figure 2: The minimum size of the editor’s bias such that the quality costs of the double-anonymous procedure outweigh its benefits (given as a percentage point difference in acceptance rates).

First, suppose that $\sigma_{in}^2 = \sigma_{sc}^2 = 1$ and $\sigma_{rv}^2 = 4$, so there is relatively little

¹⁷These are calculated without regard for whether the editor knows the author or not. In particular, the rate of acceptance for authors the editor is biased against is

$$\frac{p_{KA} \Pr(\mu_i^{KA} > q^*) + p_{UA} \Pr(\mu_i^{UA} > q^*)}{p_{KA} + p_{UA}}, \text{ and } \frac{p_{KF} \Pr(\mu_i^{KF} > q^*) + p_{UF} \Pr(\mu_i^{UF} > q^*)}{p_{KF} + p_{UF}}$$

is the rate of acceptance for authors the editor is biased in favor of.

variation in the quality of individual papers and in the average quality of authors but relatively high variation in reviewer estimates of quality. In this extreme case the triple-anonymous procedure has an acceptance rate as low as 0.72 %. If the groups are all of equal size ($p_{KA} = p_{KF} = p_{UA} = p_{UF} = 0.25$) then under the double-anonymous procedure the acceptance rate for authors the editor is biased in favor of needs to be as much as 2.66 pp higher than the acceptance rate for authors the editor is biased against, in order for the average quality under the two procedures to be equal. Clearly a 2.66 pp bias is very large for a journal that only accepts less than 1 % of papers. If the bias is any less than that there is no harm to the readers in using the double-anonymous procedure.

Second, suppose that $\sigma_{in}^2 = \sigma_{sc}^2 = 4$ and $\sigma_{rv}^2 = 1$, so the variation in quality of both papers and authors is relatively high but reviewers' estimates are relatively accurate. Then the triple-anonymous procedure has an acceptance rate of 22.66 %. If, moreover, the editor knows relatively few authors ($p_{KA} = p_{KF} = 0.15$, $p_{UA} = p_{UF} = 0.35$) then the acceptance rate for authors the editor is biased in favor of needs to be only 2.23 pp higher than the acceptance rate for authors the editor is biased against, in order for the quality costs of the double-anonymous procedure to outweigh its benefits. For a journal accepting about 23 % of papers that means that even if the identity bias of the editor is relatively mild the journal's readers are harmed if the double-anonymous procedure is used.

Based on these results, and the fact that the parameter values are unlikely to be known in practice, it is unclear whether the double-anonymous procedure or the triple-anonymous procedure will lead to a higher average quality of published papers for any particular journal.¹⁸ So in general it is not clear that an argument that the double-anonymous procedure harms the

¹⁸Note that the evidence collected by Laband and Piette (1994) does not help settle this question, as they do not directly compare the triple-anonymous and the double-anonymous procedure. Their evidence supports a positive epistemic effect of connection bias, but not a verdict on the overall epistemic effect of triple-anonymizing.

journal's readers can be made. At the same time, a general argument that the double-anonymous procedure helps the readers is not available either. Given this, I am inclined to recommend a triple-anonymous procedure for all journals because not doing so is unfair to authors.

One might be tempted to draw a different policy recommendation from this paper: use triple-anonymous review to prevent the negative effects of identity bias on quality, but provide the editor with the author's h-index or some other citation index to benefit from the reduced uncertainty associated with knowing an author's average quality. I do not endorse this suggestion for at least two reasons. First, it is unfair to authors as discussed in section 3. Second, depending on one's interpretation of quality, it may be difficult or impossible to infer author quality from citations (Lindsey 1989, Heesen forthcoming, Bright 2015).

I have argued in this section that the net effect of connection bias and identity bias on quality is unclear. But I argued in section 2 that the positive effect of connection bias may only exist in certain fields. In fields where papers rely partially on the author's testimony there is value in knowing the identity of the author. But in other fields such as mathematics and parts of the humanities testimony is not taken to play a role—the paper itself constitutes the contribution to the field—and so arguably there is no value in knowing the identity of the author.

In those fields, then, there is no quality benefit from connection bias, but there is still a quality cost from identity bias. So here the strongest case for the triple-anonymous procedure emerges, as the double-anonymous procedure is both unfair to authors and harms readers.

I have focused on evaluating triple-anonymous review, in particular in contrast to double-anonymous review. In many fields, particularly in the natural sciences, single-anonymous review is the norm, and so the more pertinent question is whether they should switch to double-anonymous review. Can the present model be used or adapted to address this question?

Analyzing a model in which both the editor and one or more reviewers display connection bias and/or identity bias is beyond the scope of this paper. Here I only remark on one relatively simple scenario: the case in which the editor does not display identity bias but the reviewer does.

Suppose the reviewer is biased against one group, reducing reviewer estimates of paper quality by ε if the author belongs to that group and raising estimates by δ otherwise. If the editor knows the reviewer is biased, she can take the reviewer's bias into account. In particular, if she knows which group the reviewer is biased against and the size of the bias, learning the biased reviewer estimate is equivalent to learning what the unbiased reviewer estimate would have been, and so a rational unbiased editor simply updates as if she learned the unbiased reviewer estimate. So in this case reviewer bias has no effect on acceptance decisions at all.

If the editor does not know the reviewer is biased, she may (naively) treat the biased reviewer estimate as an unbiased estimate. In this case the analysis is very similar to the one given in above. A close analogue of theorem 4 holds. The only difference is that the effect of the variances is flipped. High values of σ_{in}^2 and σ_{sc}^2 increase the consequences of the reviewer's bias, while high values of σ_{rv}^2 reduce it. This is the reverse of what happens in the version of the model I analyzed above (cf. proposition 11 in appendix A).

6 Conclusion

In this paper I have considered two types of arguments for triple-anonymous review: one based on fairness considerations from the perspective of the author and one based on the consequences for the readers of the journal.

I have argued that the double-anonymous procedure introduces differential treatment of scientific authors. In particular, editors are more likely to publish papers by authors they know (connection bias, theorem 1) and less likely to publish papers by authors they apply negative identity-prejudicial

stereotypes to (identity bias, theorem 4). Whenever a paper is rejected as a result of one of these biases an epistemic injustice (in the sense of Fricker 2007) is committed against the author. This is a fairness-based argument in favor of triple-anonymizing.

From the readers' perspective the story is more mixed. Generally speaking connection bias has a positive effect on the quality of published papers and identity bias a negative one. Whether the readers are better off under the triple-anonymous procedure then depends on how these effects trade off, which is highly context-dependent. This yields a more nuanced view than that suggested by either Laband and Piette (1994), who focus only on connection bias, or by the argument for triple-anonymizing based on the consequences for the readers, which focuses only on identity bias.

However, in mathematics and parts of the humanities there is arguably no positive quality effect from connection bias, as knowing about an author's other work is not taken to be relevant (Easwaran 2009). So here the negative effect of identity bias is the only relevant consideration from the readers' perspective. In this situation, considerations concerning the consequences for the author and considerations concerning the consequences for the readers point in the same direction: in favor of triple-anonymous review.

A The Acceptance Probability and the Average Quality of Papers

The following properties of the normal distribution will be useful (see, e.g., Johnson et al. 1994, chapter 13, section 3). Let $X \sim N(m, s^2)$. Then the moment-generating function of X is

$$\mathbb{E}[\exp\{tX\}] = \exp\left\{mt + \frac{1}{2}s^2t^2\right\}. \quad (1)$$

Let $Y = aX + b$ (with $a \neq 0$). Then

$$Y \sim N(am + b, a^2s^2). \quad (2)$$

Proposition 5.

$$\pi(q_i | r_i) \sim N(\mu_i^U, \sigma_{\pi|r}^2) \quad \text{and} \quad \pi(q_i | r_i, \mu_i) \sim N(\mu_i^K, \sigma_{\pi|r\mu}^2),$$

where

$$\begin{aligned} \mu_i^U &= \frac{\sigma_{in}^2 + \sigma_{sc}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2} r_i + \frac{\sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2} \mu, & \sigma_{\pi|r}^2 &= \frac{(\sigma_{in}^2 + \sigma_{sc}^2)\sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}, \\ \mu_i^K &= \frac{\sigma_{in}^2}{\sigma_{in}^2 + \sigma_{rv}^2} r_i + \frac{\sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2} \mu_i, & \sigma_{\pi|r\mu}^2 &= \frac{\sigma_{in}^2 \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}. \end{aligned}$$

See DeGroot (2004, section 9.5, or any other textbook that covers Bayesian statistics) for a proof of proposition 5.

Proposition 6. *The posterior means are normally distributed, with $\mu_i^U \sim N(\mu, \sigma_U^2)$ and $\mu_i^K \sim N(\mu, \sigma_K^2)$. Here,*

$$\sigma_U^2 = \frac{(\sigma_{in}^2 + \sigma_{sc}^2)^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2} \quad \text{and} \quad \sigma_K^2 = \frac{\sigma_{in}^4 + \sigma_{sc}^2(\sigma_{in}^2 + \sigma_{rv}^2)}{\sigma_{in}^2 + \sigma_{rv}^2}.$$

Moreover, if $\sigma_{sc}^2 > 0$ and $\sigma_{rv}^2 > 0$, then $\sigma_U^2 < \sigma_K^2$.

Proof. First consider the distribution of r_i . Since $r_i | q_i \sim N(q_i, \sigma_{rv}^2)$, $q_i | \mu_i \sim N(\mu_i, \sigma_{in}^2)$, and $\mu_i \sim N(\mu, \sigma_{sc}^2)$, it follows that $r_i | \mu_i \sim N(\mu_i, \sigma_{in}^2 + \sigma_{rv}^2)$ and $r_i \sim N(\mu, \sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2)$.

Since μ is a constant, μ_i^U is a linear transformation of r_i . By equation 2 μ_i^U is normally distributed with mean μ and variance σ_U^2 .

Determining the distribution of μ_i^K is slightly trickier because there are two random variables involved: r_i and μ_i . Define $X_i = \mu_i^K - \mu_i = \frac{\sigma_{in}^2}{\sigma_{in}^2 + \sigma_{rv}^2} (r_i - \mu_i)$. Since, $r_i | \mu_i \sim N(\mu_i, \sigma_{in}^2 + \sigma_{rv}^2)$,

$$X_i | \mu_i \sim N\left(0, \frac{\sigma_{in}^4}{\sigma_{in}^2 + \sigma_{rv}^2}\right).$$

It remains to determine the distribution of $\mu_i^K = X_i + \mu_i$. This can be done using the moment-generating function and the law of total expectation.

$$\begin{aligned}
\mathbb{E}[\exp\{t\mu_i^K\}] &= \mathbb{E}[\mathbb{E}[\exp\{tX_i + t\mu_i\} \mid \mu_i]] \\
&= \mathbb{E}[\exp\{t\mu_i\}\mathbb{E}[\exp\{tX_i\} \mid \mu_i]] \\
&= \exp\left\{0t + \frac{1}{2}\frac{\sigma_{in}^4}{\sigma_{in}^2 + \sigma_{rv}^2}t^2\right\}\mathbb{E}[\exp\{t\mu_i\}] \\
&= \exp\left\{\mu t + \frac{1}{2}\frac{\sigma_{in}^4 + \sigma_{sc}^2(\sigma_{in}^2 + \sigma_{rv}^2)}{\sigma_{in}^2 + \sigma_{rv}^2}t^2\right\},
\end{aligned}$$

which is exactly the moment-generating function of the desired normal distribution.

Finally, note that

$$\begin{aligned}
\sigma_U^2 &= \frac{(\sigma_{in}^2 + \sigma_{sc}^2)^2(\sigma_{in}^2 + \sigma_{rv}^2)}{(\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2)(\sigma_{in}^2 + \sigma_{rv}^2)}, \\
\sigma_K^2 &= \frac{(\sigma_{in}^2 + \sigma_{sc}^2)^2(\sigma_{in}^2 + \sigma_{rv}^2) + \sigma_{sc}^2\sigma_{rv}^4}{(\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2)(\sigma_{in}^2 + \sigma_{rv}^2)}.
\end{aligned}$$

So $\sigma_U^2 < \sigma_K^2$ whenever $\sigma_{sc}^2 > 0$ and $\sigma_{rv}^2 > 0$ (and $\sigma_U^2 = \sigma_K^2$ otherwise, assuming either $\sigma_{in}^2 > 0$ or $\sigma_{rv}^2 > 0$). \square

Theorem 1. $\Pr(\mu_i^K > q^*) > \Pr(\mu_i^U > q^*)$ if $q^* > \mu$, $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$.

Proof. It follows from proposition 6 that

$$\Pr(\mu_i^K > q^*) = 1 - \Phi\left(\frac{q^* - \mu}{\sigma_K}\right) \quad \text{and} \quad \Pr(\mu_i^U > q^*) = 1 - \Phi\left(\frac{q^* - \mu}{\sigma_U}\right),$$

where Φ is the distribution function (or cumulative density function) of a standard normal distribution. Since Φ is (strictly) increasing in its argument, and $\sigma_K > \sigma_U$ by proposition 6, the theorem follows immediately. \square

Proposition 7.

$$\mathbb{E}[q_i \mid \mu_i^U > q^*] = \mathbb{E}[\mu_i^U \mid \mu_i^U > q^*] \quad \text{and} \quad \mathbb{E}[q_i \mid \mu_i^K > q^*] = \mathbb{E}[\mu_i^K \mid \mu_i^K > q^*].$$

Proof. Because μ_i^U is simply an (invertible) transformation of r_i , it follows that

$$q_i \mid \mu_i^U \sim q_i \mid r_i \sim N\left(\mu_i^U, \sigma_{\pi|r}^2\right).$$

The distribution of $q_i \mid \mu_i^K$ is found using the moment-generating function and the law of total expectation:

$$\begin{aligned} \mathbb{E}[\exp\{tq_i\} \mid \mu_i^K] &= \mathbb{E}[\mathbb{E}[\exp\{tq_i\} \mid \mu_i, \mu_i^K] \mid \mu_i^K] \\ &= \mathbb{E}[\mathbb{E}[\exp\{tq_i\} \mid \mu_i, r_i] \mid \mu_i^K] \\ &= \mathbb{E}\left[\exp\left\{\mu_i^K t + \frac{1}{2}\sigma_{\pi|r\mu}^2 t^2\right\} \mid \mu_i^K\right] \\ &= \exp\left\{\mu_i^K t + \frac{1}{2}\sigma_{\pi|r\mu}^2 t^2\right\}, \end{aligned}$$

where the second equality follows because, if μ_i is given, μ_i^K is simply an invertible transformation of r_i . So:

$$q_i \mid \mu_i^K \sim q_i \mid r_i, \mu_i \sim N\left(\mu_i^K, \sigma_{\pi|r\mu}^2\right).$$

The important point here is that $\mathbb{E}[q_i \mid \mu_i^x] = \mu_i^x$ both for $x = U$ and $x = K$.

Now the law of total expectation can be used to establish that

$$\mathbb{E}[q_i \mid \mu_i^x > q^*] = \mathbb{E}[\mathbb{E}[q_i \mid \mu_i^x] \mid \mu_i^x > q^*] = \mathbb{E}[\mu_i^x \mid \mu_i^x > q^*],$$

for $x = U, K$. □

Let $X \sim N(m, s^2)$ be a normally distributed random variable. Then $X \mid X > a$ follows a *left-truncated normal distribution*, with left-truncation

point a . As a result of proposition 7 the mean of the left-truncated normal distribution is of interest. According to, e.g., Johnson et al. (1994, chapter 13, section 10.1), this mean can be expressed as

$$\mathbb{E}[X | X > a] = m + sR\left(\frac{a - m}{s}\right). \quad (3)$$

Here

$$R(x) = \frac{\phi(x)}{1 - \Phi(x)}$$

for all $x \in \mathbb{R}$, where ϕ is the probability density function of the standard normal distribution, and Φ is its distribution function. R is the inverse of what is known in the literature (e.g., Gordon 1941) as *Mills' ratio*.

It follows from the definitions that $R(x) > 0$ for all $x \in \mathbb{R}$ and that

$$R'(x) = R(x)^2 - xR(x). \quad (4)$$

Proposition 8 (Gordon (1941)). *For all $x > 0$, $R(x) < \frac{x^2+1}{x}$.*

Proposition 8 can be used to establish the next result.

Proposition 9. *If $X \sim N(m, s^2)$ and $Y \sim N(m, \sigma^2)$ with $\sigma > s > 0$ then $\mathbb{E}[Y | Y > a] > \mathbb{E}[X | X > a]$.*

Proof. It suffices to show that the derivative $\frac{\partial}{\partial s}\mathbb{E}[X | X > a]$ is positive for all $s > 0$. Differentiating equation (3) (using equation (4)) yields

$$\frac{\partial}{\partial s}\mathbb{E}[X | X > a] = \left(\left(\frac{a - m}{s}\right)^2 + 1\right)R\left(\frac{a - m}{s}\right) - \frac{a - m}{s}R\left(\frac{a - m}{s}\right)^2.$$

Since $R\left(\frac{a - m}{s}\right) > 0$, $\frac{\partial}{\partial s}\mathbb{E}[X | X > a] > 0$ if and only if

$$\left(\frac{a - m}{s}\right)^2 + 1 - \frac{a - m}{s}R\left(\frac{a - m}{s}\right) > 0.$$

This is true whenever $\frac{a-m}{s} \leq 0$ because then both terms in the sum are positive. Proposition 8 guarantees that it is true whenever $\frac{a-m}{s} > 0$ as well. \square

Theorem 2. $\mathbb{E}[q_i \mid \mu_i^K > q^*] > \mathbb{E}[q_i \mid \mu_i^U > q^*]$ whenever $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$.

Proof. By proposition 7,

$$\mathbb{E}[q_i \mid \mu_i^U > q^*] = \mathbb{E}[\mu_i^U \mid \mu_i^U > q^*] \quad \text{and} \quad \mathbb{E}[q_i \mid \mu_i^K > q^*] = \mathbb{E}[\mu_i^K \mid \mu_i^K > q^*].$$

By proposition 6, $\mu_i^U \sim N(\mu, \sigma_U^2)$ and $\mu_i^K \sim N(\mu, \sigma_K^2)$, with $\sigma_U < \sigma_K$. Hence the conditions of proposition 9 are satisfied, and the result follows. \square

Proposition 10.

$$\begin{aligned} \mu_i^K \mid \mu_i &\sim N\left(\mu_i, \frac{\sigma_{in}^4}{\sigma_{in}^2 + \sigma_{rv}^2}\right), \\ \mu_i^U \mid \mu_i &\sim N\left(\frac{\sigma_{in}^2 + \sigma_{sc}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2} \mu_i + \frac{\sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2} \mu, \frac{(\sigma_{in}^2 + \sigma_{sc}^2)(\sigma_{in}^2 + \sigma_{rv}^2)}{(\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2)^2}\right). \end{aligned}$$

Proof. Since μ_i is given and hence behaves like a constant, both μ_i^K and μ_i^U are simply linear transformations of r_i . Given that $r_i \mid \mu_i \sim N(\mu_i, \sigma_{in}^2 + \sigma_{rv}^2)$, both results follow from equation 2. \square

Theorem 3. Given $\sigma_{sc}^2 > 0$ and $\sigma_{rv}^2 > 0$,

$$\Pr(\mu_i^K > q^* \mid \mu_i) \geq \Pr(\mu_i^U > q^* \mid \mu_i) \Leftrightarrow \mu_i \geq \frac{\sigma_{in}^2}{\sigma_{in}^2 + \sigma_{sc}^2} \mu + \frac{\sigma_{sc}^2}{\sigma_{in}^2 + \sigma_{sc}^2} q^*.$$

Proof. When either $\sigma_{sc}^2 = 0$ or $\sigma_{rv}^2 = 0$, it follows from proposition 10 that $\mu_i^K \mid \mu_i \sim \mu_i^U \mid \mu_i$ and so the conditional acceptance probabilities are equal.

Otherwise,¹⁹

$$\Pr(\mu_i^K > q^* \mid \mu_i) = 1 - \Phi\left(\frac{\sqrt{\sigma_{in}^2 + \sigma_{rv}^2}}{\sigma_{in}^2}(q^* - \mu_i)\right),$$

$$\Pr(\mu_i^U > q^* \mid \mu_i) = 1 - \Phi\left(\frac{(\sigma_{in}^2 + \sigma_{sc}^2)(q^* - \mu_i) + \sigma_{rv}^2(q^* - \mu)}{(\sigma_{in}^2 + \sigma_{sc}^2)\sqrt{\sigma_{in}^2 + \sigma_{rv}^2}}\right).$$

So $\Pr(\mu_i^K > q^* \mid \mu_i) \geq \Pr(\mu_i^U > q^* \mid \mu_i)$ if and only if

$$\frac{(\sigma_{in}^2 + \sigma_{sc}^2)(q^* - \mu_i) + \sigma_{rv}^2(q^* - \mu)}{(\sigma_{in}^2 + \sigma_{sc}^2)\sqrt{\sigma_{in}^2 + \sigma_{rv}^2}} \geq \frac{\sqrt{\sigma_{in}^2 + \sigma_{rv}^2}}{\sigma_{in}^2}(q^* - \mu_i).$$

Some algebra yields the result. \square

Proposition 11.

$$\begin{aligned} \pi_A(q_i \mid r_i, \mu_i) &\sim N(\mu_i^{KA}, \sigma_{\pi|r\mu}^2), & \pi_F(q_i \mid r_i, \mu_i) &\sim N(\mu_i^{KF}, \sigma_{\pi|r\mu}^2), \\ \pi_A(q_i \mid r_i) &\sim N(\mu_i^{UA}, \sigma_{\pi|r}^2), & \pi_F(q_i \mid r_i) &\sim N(\mu_i^{UF}, \sigma_{\pi|r}^2), \end{aligned}$$

where

$$\begin{aligned} \mu_i^{KA} &= \mu_i^K - \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, & \mu_i^{KF} &= \mu_i^K + \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, \\ \mu_i^{UA} &= \mu_i^U - \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}, & \mu_i^{UF} &= \mu_i^U + \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}. \end{aligned}$$

For a proof of proposition 11 I refer once again to DeGroot (2004, section 9.5).

Proposition 12. *The posterior means are normally distributed, with*

¹⁹The expression for $\Pr(\mu_i^K > q^* \mid \mu_i)$ and the remainder of this proof assume that $\sigma_{in}^2 > 0$. If $\sigma_{in}^2 = 0$ then the desired probability is one if $\mu_i > q^*$ and zero otherwise. Since $0 < \Pr(\mu_i^U > q^* \mid \mu_i) < 1$, the result follows.

$$\begin{aligned}
\mu_i^{KA} &\sim N\left(\mu - \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, \sigma_K^2\right), \\
\mu_i^{KF} &\sim N\left(\mu + \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, \sigma_K^2\right), \\
\mu_i^{UA} &\sim N\left(\mu - \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}, \sigma_U^2\right), \\
\mu_i^{UF} &\sim N\left(\mu + \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}, \sigma_U^2\right).
\end{aligned}$$

Proof. Since μ_i^{KA} and μ_i^{KF} are simply μ_i^K shifted by a constant (see proposition 11) they follow the same distribution as μ_i^K except that its mean is shifted by the same constant. Similarly μ_i^{UA} and μ_i^{UF} are just μ_i^U shifted by a constant. So the results follow from proposition 6. \square

For notational convenience, I introduce q^{KA} , q^{KF} , q^{UA} , and q^{UF} , defined by

$$\begin{aligned}
q^{KA} &= q^* + \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, & q^{KF} &= q^* - \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{rv}^2}, \\
q^{UA} &= q^* + \frac{\varepsilon \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}, & q^{UF} &= q^* - \frac{\delta \cdot \sigma_{rv}^2}{\sigma_{in}^2 + \sigma_{sc}^2 + \sigma_{rv}^2}.
\end{aligned}$$

Theorem 4. *If $\varepsilon + \delta > 0$, $\sigma_{sc}^2 > 0$, and $\sigma_{rv}^2 > 0$,*

$$\Pr(\mu_i^{KA} > q^*) < \Pr(\mu_i^{KF} > q^*) \quad \text{and} \quad \Pr(\mu_i^{UA} > q^*) < \Pr(\mu_i^{UF} > q^*).$$

Proof. For the first inequality, note that

$$\Pr(\mu_i^{KA} > q^*) = 1 - \Phi\left(\frac{q^{KA} - \mu}{\sigma_K}\right) < 1 - \Phi\left(\frac{q^{KF} - \mu}{\sigma_K}\right) = \Pr(\mu_i^{KF} > q^*).$$

The equalities follow from the distributions of the posterior means established in proposition 12. The inequality follows from the fact that Φ is strictly increasing in its argument. By the same reasoning,

$$\Pr(\mu_i^{UA} > q^*) = 1 - \Phi\left(\frac{q^{UA} - \mu}{\sigma_U}\right) < 1 - \Phi\left(\frac{q^{UF} - \mu}{\sigma_U}\right) = \Pr(\mu_i^{UF} > q^*). \quad \square$$

Proposition 13.

$$\begin{aligned} \Pr(A_i) &= p_{KA} \left(1 - \Phi\left(\frac{q^{KA} - \mu}{\sigma_K}\right)\right) + p_{KF} \left(1 - \Phi\left(\frac{q^{KF} - \mu}{\sigma_K}\right)\right) \\ &\quad + p_{UA} \left(1 - \Phi\left(\frac{q^{UA} - \mu}{\sigma_U}\right)\right) + p_{UF} \left(1 - \Phi\left(\frac{q^{UF} - \mu}{\sigma_U}\right)\right). \\ \mathbb{E}[q_i | A_i] &= \mu + \frac{\sigma_K}{\Pr(A_i)} \left(p_{KA} \phi\left(\frac{q^{KA} - \mu}{\sigma_K}\right) + p_{KF} \phi\left(\frac{q^{KF} - \mu}{\sigma_K}\right)\right) \\ &\quad + \frac{\sigma_U}{\Pr(A_i)} \left(p_{UA} \phi\left(\frac{q^{UA} - \mu}{\sigma_U}\right) + p_{UF} \phi\left(\frac{q^{UF} - \mu}{\sigma_U}\right)\right). \end{aligned}$$

Proof. The expression for $\Pr(A_i)$ follows immediately from the distributions of the posterior means established in proposition 12.

To get an expression for $\mathbb{E}[q_i | A_i]$, consider first the average quality of scientist i 's paper given that it is accepted and given that scientist i is in the group of scientists known to the editor that the editor is biased against. This average quality is

$$\mathbb{E}[q_i | \mu_i^{KA} > q^*] = \mathbb{E}[\mu_i^K | \mu_i^K > q^{KA}] = \mu + \sigma_K R\left(\frac{q^{KA} - \mu}{\sigma_K}\right),$$

where the first equality uses the fact that $\mu_i^{KA} > q^*$ is equivalent to $\mu_i^K > q^{KA}$ and then applies proposition 7, and the second equality uses equation 3. Similarly,

$$\begin{aligned}\mathbb{E} [q_i | \mu_i^{KF} > q^*] &= \mu + \sigma_K R \left(\frac{q^{KF} - \mu}{\sigma_K} \right), \\ \mathbb{E} [q_i | \mu_i^{UA} > q^*] &= \mu + \sigma_U R \left(\frac{q^{UA} - \mu}{\sigma_U} \right), \\ \mathbb{E} [q_i | \mu_i^{UF} > q^*] &= \mu + \sigma_U R \left(\frac{q^{UF} - \mu}{\sigma_U} \right).\end{aligned}$$

The average quality of accepted papers $\mathbb{E}[q_i | A_i]$ is a weighted sum of these expectations. The weights are given by the proportion of accepted papers that are written by a scientist in that particular group. For example, authors known to the editor that she is biased against form a $p_{KA} \Pr(\mu_i^{KA} > q^*) / \Pr(A_i)$ proportion of accepted papers. Hence

$$\begin{aligned}\mathbb{E} [q_i | A_i] &= \frac{1}{\Pr(A_i)} p_{KA} \Pr(\mu_i^{KA} > q^*) \mathbb{E} [q_i | \mu_i^{KA} > q^*] \\ &\quad + \frac{1}{\Pr(A_i)} p_{KF} \Pr(\mu_i^{KF} > q^*) \mathbb{E} [q_i | \mu_i^{KF} > q^*] \\ &\quad + \frac{1}{\Pr(A_i)} p_{UA} \Pr(\mu_i^{UA} > q^*) \mathbb{E} [q_i | \mu_i^{UA} > q^*] \\ &\quad + \frac{1}{\Pr(A_i)} p_{UF} \Pr(\mu_i^{UF} > q^*) \mathbb{E} [q_i | \mu_i^{UF} > q^*] \\ &= \mu + \frac{\sigma_K}{\Pr(A_i)} \left(p_{KA} \phi \left(\frac{q^{KA} - \mu}{\sigma_K} \right) + p_{KF} \phi \left(\frac{q^{KF} - \mu}{\sigma_K} \right) \right) \\ &\quad + \frac{\sigma_U}{\Pr(A_i)} \left(p_{UA} \phi \left(\frac{q^{UA} - \mu}{\sigma_U} \right) + p_{UF} \phi \left(\frac{q^{UF} - \mu}{\sigma_U} \right) \right). \quad \square\end{aligned}$$

References

Charles D. Bailey, Dana R. Hermanson, and Timothy J. Louwers. An examination of the peer review process in accounting journals. *Jour-*

- nal of Accounting Education*, 26(2):55–72, 2008a. ISSN 0748-5751. doi: 10.1016/j.jaccedu.2008.04.001. URL <http://www.sciencedirect.com/science/article/pii/S0748575108000201>.
- Charles D. Bailey, Dana R. Hermanson, and James G. Tompkins. The peer review process in finance journals. *Journal of Financial Education*, 34: 1–27, 2008b. ISSN 0093-3961. URL <http://www.jstor.org/stable/41948838>.
- Damien Besancenot, Kim V. Huynh, and Joao R. Faria. Search and research: the influence of editorial boards on journals’ quality. *Theory and Decision*, 73(4):687–702, 2012. ISSN 0040-5833. doi: 10.1007/s11238-012-9314-7. URL <http://dx.doi.org/10.1007/s11238-012-9314-7>.
- Rebecca M. Blank. The effects of double-blind versus single-blind reviewing: Experimental evidence from The American Economic Review. *The American Economic Review*, 81(5):1041–1067, 1991. ISSN 00028282. URL <http://www.jstor.org/stable/2006906>.
- Denny Borsboom, Jan-Willem Romeijn, and Jelte M. Wicherts. Measurement invariance versus selection invariance: Is fair selection possible? *Psychological Methods*, 13(2):75–98, Jun 2008. doi: 10.1037/1082-989X.13.2.75. URL <http://dx.doi.org/10.1037/1082-989X.13.2.75>.
- Liam Kofi Bright. Against candidate quality. Manuscript, 2015. URL https://www.academia.edu/11673059/Against_Candidate_Quality.
- Liam Kofi Bright, Daniel Malinsky, and Morgan Thompson. Causally interpreting intersectionality theory. *Philosophy of Science*, 83(1):60–81, 2016. ISSN 00318248. URL <http://www.jstor.org/stable/10.1086/684173>.
- Justin Bruner and Cailin O’Connor. Power, bargaining, and collaboration. In Thomas Boyer-Kassem, Conor Mayo-Wilson, and Michael Weisberg,

- editors, *Scientific Collaboration and Collective Knowledge*. Oxford University Press, Oxford, forthcoming. URL <http://philpapers.org/rec/BRUPBA-2>.
- Amber E. Budden, Tom Tregenza, Lonnie W. Aarssen, Julia Koricheva, Roosa Leimu, and Christopher J. Lortie. Double-blind review favours increased representation of female authors. *Trends in Ecology & Evolution*, 23(1):4–6, 2008. ISSN 0169-5347. doi: 10.1016/j.tree.2007.07.008. URL <http://www.sciencedirect.com/science/article/pii/S0169534707002704>.
- Kenneth E. Clark. *America's Psychologists: A Survey of a Growing Profession*. American Psychological Association, Washington, 1957.
- Jonathan R. Cole and Stephen Cole. Measuring the quality of sociological research: Problems in the use of the “Science Citation Index”. *The American Sociologist*, 6(1):23–29, 1971. ISSN 00031232. URL <http://www.jstor.org/stable/27701705>.
- Stephen Cole and Jonathan R. Cole. Scientific output and recognition: A study in the operation of the reward system in science. *American Sociological Review*, 32(3):377–390, 1967. ISSN 00031224. URL <http://www.jstor.org/stable/2091085>.
- Stephen Cole and Jonathan R. Cole. Visibility and the structural bases of awareness of scientific research. *American Sociological Review*, 33(3):397–413, 1968. ISSN 00031224. URL <http://www.jstor.org/stable/2091914>.
- Patricia Hill Collins and Valerie Chepp. Intersectionality. In Georgina Waylen, Karen Celis, Johanna Kantola, and S. Laurel Weldon, editors, *The Oxford Handbook of Gender and Politics*, chapter 2, pages 57–87. Oxford University Press, Oxford, 2013. ISBN 0199751455.

- Rick Crandall. Editorial responsibilities in manuscript review. *Behavioral and Brain Sciences*, 5:207–208, Jun 1982. ISSN 1469-1825. doi: 10.1017/S0140525X00011316. URL http://journals.cambridge.org/article_S0140525X00011316.
- Diana Crane. The gatekeepers of science: Some factors affecting the selection of articles for scientific journals. *The American Sociologist*, 2(4):195–201, 1967. ISSN 00031232. URL <http://www.jstor.org/stable/27701277>.
- Morris H. DeGroot. *Optimal Statistical Decisions*. John Wiley & Sons, New Jersey, 2004.
- Kenny Easwaran. Probabilistic proofs and transferability. *Philosophia Mathematica*, 17(3):341–362, 2009. doi: 10.1093/philmat/nkn032. URL <http://philmat.oxfordjournals.org/content/17/3/341.abstract>.
- Glenn Ellison. Evolving standards for academic publishing: A q-r theory. *Journal of Political Economy*, 110(5):994–1034, 2002. ISSN 00223808. URL <http://www.jstor.org/stable/10.1086/341871>.
- João Ricardo Faria. The game academics play: Editors versus authors. *Bulletin of Economic Research*, 57(1):1–12, 2005. ISSN 1467-8586. doi: 10.1111/j.1467-8586.2005.00212.x. URL <http://dx.doi.org/10.1111/j.1467-8586.2005.00212.x>.
- Miranda Fricker. *Epistemic Injustice: Power and the Ethics of Knowing*. Oxford University Press, Oxford, 2007.
- Robert D. Gordon. Values of Mills' ratio of area to bounding ordinate and of the normal probability integral for large values of the argument. *The Annals of Mathematical Statistics*, 12(3):364–366, 1941. ISSN 00034851. URL <http://www.jstor.org/stable/2235868>.

- Remco Heesen. Academic superstars: Competent or lucky? *Synthese*, forthcoming. ISSN 1573-0964. doi: 10.1007/s11229-016-1146-5. URL <http://dx.doi.org/10.1007/s11229-016-1146-5>.
- David L. Hull. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. University of Chicago Press, Chicago, 1988. ISBN 0226360504.
- Norman L. Johnson, Samuel Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 1. John Wiley & Sons, New York, second edition, 1994.
- Philip Kitcher. *The Advancement of Science: Science without Legend, Objectivity without Illusions*. Oxford University Press, Oxford, 1993. ISBN 0195046285.
- David N. Laband. Publishing favoritism: A critique of department rankings based on quantitative publishing performance. *Southern Economic Journal*, 52(2):510–515, 1985. ISSN 00384038. URL <http://www.jstor.org/stable/1059636>.
- David N. Laband and Michael J. Piette. Favoritism versus search for good papers: Empirical evidence regarding the behavior of journal editors. *Journal of Political Economy*, 102(1):194–203, 1994. ISSN 00223808. URL <http://www.jstor.org/stable/2138799>.
- Carole J. Lee and Christian D. Schunn. Philosophy journal practices and opportunities for bias. *American Philosophical Association Newsletter on Feminism and Philosophy*, 10(1):5–10, 2010. ISSN 2155-9708. URL http://www.apaonline.org/?feminism_newsletter.
- Carole J. Lee, Cassidy R. Sugimoto, Guo Zhang, and Blaise Cronin. Bias in peer review. *Journal of the American Society for Information Science and*

- Technology*, 64(1):2–17, 2013. ISSN 1532-2890. doi: 10.1002/asi.22784. URL <http://dx.doi.org/10.1002/asi.22784>.
- D. Lindsey. Using citation counts as a measure of quality in science: Measuring what’s measurable rather than what’s valid. *Scientometrics*, 15(3–4):189–203, 1989. ISSN 0138-9130. doi: 10.1007/BF02017198. URL <http://dx.doi.org/10.1007/BF02017198>.
- Christopher D. Mackie. *Canonizing Economic Theory: How Theories and Ideas Are Selected in Economics*. M. E. Sharpe, New York, 1998. ISBN 9780765602848.
- Conor Mayo-Wilson, Kevin J. S. Zollman, and David Danks. The independence thesis: When individual and social epistemology diverge. *Philosophy of Science*, 78(4):653–677, 2011. ISSN 00318248. URL <http://www.jstor.org/stable/10.1086/661777>.
- Marshall H. Medoff. Editorial favoritism in economics? *Southern Economic Journal*, 70(2):425–434, 2003. ISSN 00384038. URL <http://www.jstor.org/stable/3648979>.
- Robert K. Merton. A note on science and democracy. *Journal of Legal and Political Sociology*, 1(1–2):115–126, 1942.
- Edward M. Miller. The relevance of group membership for personnel selection: A demonstration using Bayes’ theorem. *The Journal of Social, Political, and Economic Studies*, 19(3):323–358, 1994. ISSN 0278-839X.
- Corinne A. Moss-Racusin, John F. Dovidio, Victoria L. Brescoll, Mark J. Graham, and Jo Handelsman. Science faculty’s subtle gender biases favor male students. *Proceedings of the National Academy of Sciences*, 109(41):16474–16479, 2012. doi: 10.1073/pnas.1211286109. URL <http://www.pnas.org/content/109/41/16474.abstract>.

- Michael J. Piette and Kevin L. Ross. A study of the publication of scholarly output in economics journals. *Eastern Economic Journal*, 18(4):429–436, 1992. ISSN 00945056. URL <http://www.jstor.org/stable/40325474>.
- Jennifer Saul. Implicit bias, stereotype threat, and women in philosophy. In Katrina Hutchison and Fiona Jenkins, editors, *Women in Philosophy: What Needs to Change?*, chapter 2, pages 39–60. Oxford University Press, Oxford, 2013.
- Daniel L. Sherrell, Joseph F. Hair, Jr., and Mitch Griffin. Marketing academicians' perceptions of ethical research and publishing behavior. *Journal of the Academy of Marketing Science*, 17(4):315–324, 1989. ISSN 0092-0703. doi: 10.1007/BF02726642. URL <http://dx.doi.org/10.1007/BF02726642>.
- Kenneth J. Smith and Robert F. Dombrowski. An examination of the relationship between author-editor connections and subsequent citations of auditing research articles. *Journal of Accounting Education*, 16(3–4):497–506, 1998. ISSN 0748-5751. doi: 10.1016/S0748-5751(98)00019-0. URL <http://www.sciencedirect.com/science/article/pii/S0748575198000190>.
- Richard Snodgrass. Single- versus double-blind reviewing: An analysis of the literature. *ACM SIGMOD Record*, 35(3):8–21, Sep 2006. ISSN 0163-5808. doi: 10.1145/1168092.1168094. URL <http://doi.acm.org/10.1145/1168092.1168094>.
- Rhea E. Steinpreis, Katie A. Anders, and Dawn Ritzke. The impact of gender on the review of the curricula vitae of job applicants and tenure candidates: A national empirical study. *Sex Roles*, 41(7–8):509–528, 1999. ISSN 0360-0025. doi: 10.1023/A:1018839203698. URL <http://dx.doi.org/10.1023/A:1018839203698>.

- Brandon D. Stewart and B. Keith Payne. Bringing automatic stereotyping under control: Implementation intentions as efficient means of thought control. *Personality and Social Psychology Bulletin*, 34(10):1332–1345, 2008. doi: 10.1177/0146167208321269. URL <http://psp.sagepub.com/content/34/10/1332.abstract>.
- Michael Strevens. The role of the priority rule in science. *The Journal of Philosophy*, 100(2):55–79, 2003. ISSN 0022362X. URL <http://www.jstor.org/stable/3655792>.
- Shelley Tremain. *Foucault and Feminist Philosophy of Disability*. University of Michigan Press, Ann Arbor, 2017.
- Eric Luis Uhlmann and Geoffrey L. Cohen. “I think it, therefore it’s true”: Effects of self-perceived objectivity on hiring discrimination. *Organizational Behavior and Human Decision Processes*, 104(2):207–223, 2007. ISSN 0749-5978. doi: 10.1016/j.obhdp.2007.07.001. URL <http://www.sciencedirect.com/science/article/pii/S0749597807000611>.
- Virginia Valian. *Why So Slow? The Advancement of Women*. MIT Press, Cambridge, 1999. ISBN 9780262720311.
- Christine Wennerås and Agnes Wold. Nepotism and sexism in peer-review. *Nature*, 387(6631):341–343, May 1997. ISSN 0028-0836. doi: 10.1038/387341a0. URL <http://dx.doi.org/10.1038/387341a0>.
- Alan J. Ziobrowski and Karen M. Gibler. Factors academic real estate authors consider when choosing where to submit a manuscript for publication. *Journal of Real Estate Practice and Education*, 3(1):43–54, 2000. ISSN 1521-4842. URL <http://ares.metapress.com/content/1762151051KM2227>.
- Kevin J. S. Zollman. Optimal publishing strategies. *Episteme*, 6:185–199,

Jun 2009. ISSN 1750-0117. doi: 10.3366/E174236000900063X. URL http://journals.cambridge.org/article_S174236000001283.