

ON EQUITABLE NON-ANONYMOUS REVIEW

ABSTRACT. Remco Heesen has recently argued in favor of the editorial practice of triple-anonymous review on the grounds that “an injustice is committed against certain authors” under non-anonymous review. On the other hand, he concedes that the information waste of triple-anonymous review does handicap editors, in particular sacrificing a boost in the average quality of accepted papers that would otherwise be conferred by non-anonymous review. In this paper it is observed that by devoting comparatively greater reviewing resources to the papers of unfamiliar authors, editors practicing non-anonymous review can, without loss of information, avoid subjecting authors to the sorts of injustices observed by Heesen. Thus they can reap the efficiency gains of non-anonymous review without sacrificing fairness.

1. INTRODUCTION

Remco Heesen (2018) invites us to consider a simplified scholarly¹ community having a single journal. Scholars submit papers (one each, chosen at random from their output) to the journal’s editor. Some scholars tend to produce better papers than others. Let μ_i be the mean quality of a paper produced by scholar i , and assume that $\mu_i \sim N(0, 1)$. That is, if a scholar is chosen at random, then the mean quality of their papers lies in a standard normal distribution.² Papers by a fixed scholar meanwhile are assumed to lie in a normal distribution of variance 1. So, if q_i is a random paper produced by scholar i then $q_i \sim N(\mu_i, 1)$. The editor sends each submitted paper out for review, obtaining from the reviewers an aggregate estimate r_i of the quality q_i of the paper. Suppose that $r_i \sim N(q_i, 1)$. (In particular, the score is unbiased; its expectation is the actual quality of the paper.) Assume that there is a threshold $q^* > 0$ such that the editor accepts a submitted paper just when her posterior expectation of its quality is above q^* ; for simplicity I’ll take $q^* = 2$.

Now it happens³ that $E(q_i|r_i) = \frac{2}{3}r_i$. Indeed, this is a consequence of the following theorem. (Take $W = q_i$, so that $\mu_W = 0$ and $\sigma_W^2 = 2$, and $Z = r_i - q_i$, so that $\sigma_Z^2 = 1$ and $x = W + Z = r_i$.)

¹Heesen speaks of “scientific” communities rather than “scholarly” communities, but the issues he addresses are relevant more generally.

²I am assuming a variance of 1 for each of the random elements that affect the reported value r_i ; Heesen makes no such assumption, but I see no reason not to, as it simplifies the presentation without sacrificing any of the qualitative features that make the model interesting.

³An odd feature of this model is that the reviewer herself (should she know the relevant variances) would, post-review, put the expectation of the paper’s quality at $\frac{2}{3}r_i$ —even though it was she that estimated the quality of the paper to be r_i ! So there is a supposition here that reviewers are not making this sort of post-mortem “regression to the mean” correction prior to submitting their scores. (I don’t find this supposition realistic; real-life reviewers are conservative; tending in particular to underestimate the quality of strong papers.) This observation can be glossed over, though, for if reviewers are doing the post-mortem prior to report and the editor knows this then she can reconstruct the original “raw” estimate and proceed as in the text.

Theorem 1. *Suppose that $W \sim N(\mu_W, \sigma_W^2)$ and $Z \sim N(0, \sigma_Z^2)$ are independent. Then conditional on $W + Z = x$, $W \sim N(\frac{\sigma_Z^2}{\sigma_Z^2 + \sigma_W^2} \mu_W + \frac{\sigma_W^2}{\sigma_Z^2 + \sigma_W^2} x, \frac{\sigma_W^2 \sigma_Z^2}{\sigma_Z^2 + \sigma_W^2})$.*

So if the editor's only evidence is the value r_i , she will accept the paper if and only if $r_i > 3$; since $\text{Var}(r_i) = 3$, the probability of acceptance is therefore $Z(\sqrt{3}) \approx .0416$.

But the key feature of Heesen's model is that for some of the papers, the editor *does* have additional evidence—she is acquainted with the authors of these papers, in the sense that she knows, for these authors i , the corresponding values μ_i . (She knows, for these authors, the average quality of their papers.) Again invoking Theorem 1 (with $W = q_i | \mu_i$, so that $\mu_W = \mu_i$ and $\sigma_W^2 = 1$, and $Z = r_i - q_i$ as before), $E(q_i | \mu_i, r_i) = \frac{\mu_i + r_i}{2}$. So for the group the editor will accept paper i if and only if $\frac{\mu_i + r_i}{2} > 2$. Since $\text{Var}(\mu_i + r_i) = \text{Var}(2\mu_i) + \text{Var}(q_i - \mu_i) + \text{Var}(r_i - q_i) = 4 + 1 + 1 = 6$, the probability of acceptance is therefore $Z(\frac{4}{\sqrt{6}}) \approx .0512$.

Heesen concludes: “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.”

I contend that this result is not surprising. For conditionalization on any partition whatsoever would increase (strictly, in all but the extreme case of a partition independent of the conditional expectation of the quality of the paper with respect to the algebra generated by the existing review) the variance of the editor's posterior expectation of the paper's quality.⁴ And while this does not guarantee an increase in the probability of acceptance in general,⁵ it does guarantee this if the editor's posterior expectation of the paper's quality is itself assumed to have a normal distribution. (As is the case in Heesen's model.) So the phenomenon on display here is more general. What causes the acceptance rate in one of the two pools of papers to be higher is simply that the editor has more information about the papers in that pool.

Heesen opines: “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 (scholars) 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 (scholar) who submits a paper? This may well be too much to ask.” So although he ultimately concludes that the editor is “ethically and epistemically culpable” for the injustice that some scholars suffer under this system and implicitly recommends that editors should correct this injustice, his language suggests that the editor's role in the injustice is passive.

⁴The general claim is that if X is a random variable on probability space (Ω, μ) with $E(X) = 0$ and \mathcal{A} is a finer σ -algebra than \mathcal{B} then $\int E(X|\mathcal{A})^2 d\mu \geq \int E(X|\mathcal{B})^2 d\mu$. For $RHS = \int (\int X d\mu_y)^2 d\mu(y) = \int (\int \int X d\mu_{yz} d\mu_y(z))^2 d\mu(y) \leq \int \int (\int X d\mu_{yz})^2 d\mu_y(z) d\mu(y) = LHS$. (Here (μ_y) is the decomposition of μ over \mathcal{B} and (μ_{yz}) is the decomposition of μ_y over \mathcal{A} .)

⁵Here is a counterexample. Imagine that the editor learns the truth or falsity of the event $q_i > 1.9$, and nothing else. Then she will accept the paper if and only if $q_i > 1.9$. But if she learned the true quality, she would accept the paper if and only if $q_i > 2$. So here learning more yields a lower probability of acceptance.

Against this, I would say that the editor’s role is active. To see why, assume for the sake of convenience that in the original model two reviewer scores were solicited and averaged in the computation of r_i . (The variance of each of these scores is 2 conditional on q_i and they are independent conditional on q_i , so their average, r_i , has variance 1 conditional on q_i , as described above.) Then she could have achieved fairness by sending papers written by familiar authors out for review only once. For in that case Theorem 1 (with $W = q_i|\mu_i$, so that $\mu_W = \mu_i$ and $\sigma_W^2 = 1$, and $Z = r_i - q_i$, so that $\sigma_Z^2 = 2$) would yield a posterior expectation of paper quality equal to $\frac{2\mu_i+r_i}{3} = \mu_i + \frac{1}{3}(r_i - \mu_i)$. So the editor would accept the paper if and only if $\mu_i + \frac{1}{3}(r_i - \mu_i) > 2$. Since $\text{Var}(\mu_i + \frac{1}{3}(r_i - \mu_i)) = \frac{4}{3}$, the implied acceptance rate is $Z(\sqrt{3}) \approx .0416$, the same as the acceptance rate in the group of unfamiliar authors.⁶

What makes the editor’s role active, then, is that she chose to send the papers by the familiar authors out for review a second time. Yes, it “just so happened” that she started out having more information about papers by scholars that she knows. But it didn’t “just so happen” that things ended up that way. That was her doing.

Concomitantly I disagree with Heesen’s recommendation, which is elucidated in: “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 is prevented from using information about (scholars) she knows in her evaluation.” I don’t dispute that this levels the playing field (in theory), but the cost in wasted information (a debt that must ultimately be cashed out in either a decrease in quality of accepted papers or the additional sweat of reviewers) is prohibitive. A more efficient way to equalize the prospects of familiar and unfamiliar authors is to expend comparatively fewer reviewing resources on papers by familiar authors.

This argument doesn’t require that the editor is the oracle that Heesen’s model makes her out to be. Its recommendations stand (with weaker effect) in a case where the editor’s estimates of the μ_i are (normally distributed and) unbiased. Since few real-life editors approximate even this premise, the argument doesn’t advocate strongly for non-anonymous review in practice. I can’t resist mentioning however that mathematics uses non-anonymous review and it seems to work extremely well, whereas philosophy uses anonymous review and it seems to work extremely badly. Laying aside more dispiriting hypotheses as to why, Heesen’s work indicates a purely statistical one...even when review is officially triple anonymous, editors will still know some authors’ identities—comparatively more often for “well-connected” authors. The official reason we give for shopping a paper at talks prior to submission is that such vetting might improve the papers. Heesen’s work indicates a more cynical motive, however...subversion of the “anonymity” of review. This suggests that (even double) anonymous review may act more as cover for connection bias than as solution to it.

References

Heesen, Remco. 2018. When journal editors play favorites. *Philosophical Studies* 175:831-858.

⁶It may also be worth noting that the posterior variances of the q_i are equal (they are both $\frac{2}{3}$) in the two groups, so the expected quality of an accepted paper will be independent of familiarity.