

On Testing for Biases in Peer Review

Ivan Stelmakh, Nihar B. Shah and Aarti Singh

School of Computer Science
Carnegie Mellon University
{stiv,nihars,aarti}@cs.cmu.edu

Abstract

We consider the issue of biases in scholarly research, specifically, in peer review. There is a long standing debate on whether exposing author identities to reviewers induces biases against certain groups, and our focus is on designing tests to detect the presence of such biases. Our starting point is a remarkable recent work by Tomkins, Zhang and Heavlin which conducted a controlled, large-scale experiment to investigate existence of biases in the peer reviewing of the WSDM conference. We present two sets of results in this paper. The first set of results is negative, and pertains to the statistical tests and the experimental setup used in the work of Tomkins et al. We show that the test employed therein does not guarantee control over false alarm probability and under correlations between relevant variables coupled with any of the following conditions, with high probability, can declare a presence of bias when it is in fact absent: (a) measurement error, (b) model mismatch, (c) reviewer calibration. Moreover, we show that the setup of their experiment may itself inflate false alarm probability if (d) bidding is performed in non-blind manner or (e) popular reviewer assignment procedure is employed. Our second set of results is positive and is built around a novel approach to testing for biases that we propose. We present a general framework for testing for biases in (single vs. double blind) peer review. We then design hypothesis tests that under minimal assumptions guarantee control over false alarm probability and non-trivial power even under conditions (a)–(c) as well as propose an alternative experimental setup which mitigates issues (d) and (e). Finally, we show that no statistical test can improve over the non-parametric tests we consider in terms of the assumptions required to control for the false alarm probability.

1 Introduction

Past research in social sciences indicates that humans display various biases including gender, race and age biases in many critical domains such as hiring (Bertrand and Mullainathan, 2004), university admission (Thornhill, 2018), bail decisions (Arnold et al., 2018) and many others. Our focus is on fairness in academia and scholarly research, and specifically, on biases in peer review. Peer review is a backbone of scholarly research and is employed by a vast majority of journals and conferences. Due to the widespread prevalence of the Matthew effect – rich get richer and poor get poorer – in academia (Thorngate and Chowdhury, 2014; Squazzoni and Gandelli, 2012), any biases in peer review can have far reaching consequences on career trajectories of researchers. Specifically, we follow the long-standing debate (Blank, 1991; Seeber and Bacchelli, 2017; Snodgrass, 2006; Largent and Snodgrass, 2016; Okike et al., 2016; Budden et al., 2008; Webb et al., 2008; Hill and J. Provost, 2003, and references therein) on whether the authors’ identities should be hidden from reviewers or not. *The focus of this paper is on designing statistical tests to detect the presence of biases in peer review.*

In a recent remarkable piece of work, Tomkins et al. (2017) conducted a large scale (semi-) randomized controlled trial during the peer review for the ACM International Conference on Web Search and Data Mining (WSDM) 2017. In their experiment, the entire pool of reviewers was partitioned uniformly at random into

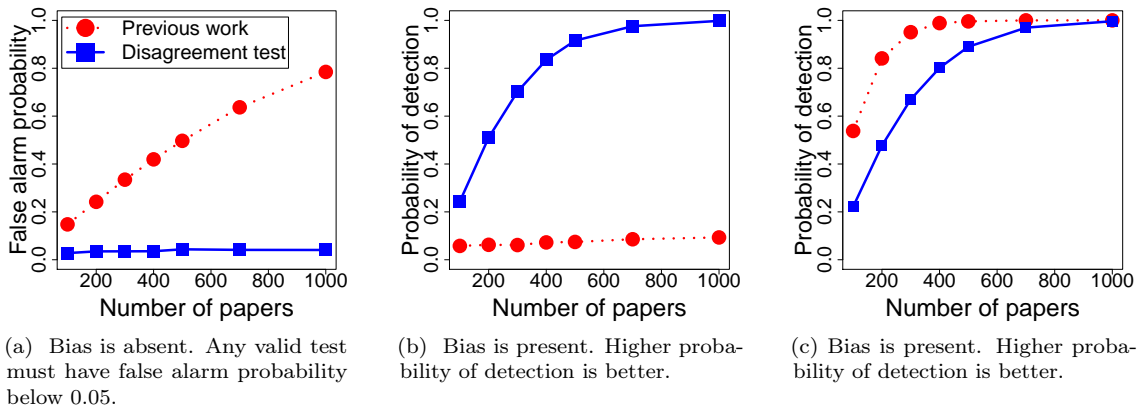


Figure 1: Synthetic simulations evaluating performance of the test in Tomkins et al. (2017) (“previous work”) and the test proposed in this paper (“DISAGREEMENT test”). Subfigures (a) and (b) are in presence of correlations and noisy estimates of true scores by double-blind reviewers; subfigure (c) has zero correlations and perfect estimate of true scores by double-blind reviewers. Details of the simulation setup are provided in Section 3. The error bars are too small to be visible.

two equal groups – single blind and double blind – and each paper was assigned to two reviewers from each of the groups. In this manner, the peer-review data contained both single-blind and double-blind reviews for each paper. The experiment allowed them to conduct a causal inference to test for biases, and conclude that the single-blind system induces a bias in favor of papers authored by (i) researchers from top-universities, (ii) researchers from top companies and (iii) famous authors. Interestingly, no bias against female-authored submissions was detected by their test, though a meta-analysis confirmed the presence of such bias. The conclusions of this experiment have had a significant impact. For instance, the WSDM conference itself completely switched to double-blind peer review starting 2018.

Testing for the presence of hypothesized phenomena is a common task in various branches of science including the biological, social, and physical sciences. The general approach therein is to impose a hard constraint on the probability of false alarm (claiming existence of the phenomenon when there is none; also called Type-I error) to some predefined threshold called significance level typically set as 0.05 or 0.01. The test would then aim to maximize the probability of detecting the phenomenon when it is actually present, while not violating the aforementioned hard constraint. The present paper also follows this general approach, for the specific setting of testing for biases using single versus double blind reviewing.

Contributions. In this paper, we study the problem of detecting bias in peer review, and present two sets of results.

(1) Analysis of prior work (Section 3) We first analyze the testing procedure used by Tomkins et al. (2017), and show that under plausible conditions the statistical test employed therein does not control for false alarm probability. In other words, *we show that under reasonable conditions, the test used by Tomkins et al. (2017) can, with probability as large as 0.5 or higher, declare the presence of a bias when the bias is in fact absent* (even when the test is tuned to have a false alarm error rate below 0.05). Specifically, we show that in presence of correlations that are reasonable to expect, any of the following factors breaks their false alarm probability guarantees: (a) measurement error caused by noise or subjectivity of reviewers, (b) model mismatch caused by violation of strong parametric assumptions on reviewers’ behavior and (c) reviewer’s calibration if she/he reviews more than one paper. Figures 1a and 1b illustrate the effect of measurement error on the false alarm probability and probability of detection of the test used by Tomkins et al. The issues we identify suggest that their test is at risk of committing Type-I error in declaring biases in their analysis.

Moving beyond the specific test used in Tomkins et al. (2017), we also study the effect of their experimental design, which is simply the standard peer-review procedure with an additional random partition of reviewers into single and double blind groups. We show that two factors – (d) asymmetrical bidding procedure and (e) non-random assignment of papers to referees – as is common in peer-review procedures today may introduce spurious correlations in the data, breaking some key independence assumptions and thereby violating the requisite guarantees on testing.

(2) Novel approach to testing for biases (Sections 4 - 6) We propose a general framework for the design of statistical tests to detect biases in this problem setting, that overcomes the aforementioned limitations. Specifically, our framework does not assume objectivity of reviewers and does not make any parametric assumptions on reviewers’ behaviour. Conceptually, we propose to think of this problem as an instance of a two-sample testing problem where single-blind and double-blind reviews form two samples and the test operates on these samples. (In contrast, Tomkins et al. (2017) study the problem under one-sample testing paradigm, operating on reviews of single-blind reviewers and using double-blind reviews to estimate some parameters in their parametric model).

We then design computationally-efficient hypothesis testing procedures that under minimal assumptions guarantee a provable control over the false alarm probability under various conditions, including aforementioned conditions (a) - (c). We supplement these tests with an alternative design of the experimental setup which coupled with our tests mitigates issues (d) - (e) while not restricting the choice of assignment algorithm.

Our tests also have non-trivial power in that they have considerably higher probability of detection in hard cases where test used by Tomkins et al. fails, and a power comparable to that of Tomkins et al. when their assumptions are exactly met. The performance of one of these tests is illustrated in Figure 1. Additionally, we show that assumptions required by our tests to control for the Type-I error rate are essentially minimal in that they cannot be further relaxed without making reliable testing impossible.

It is important to note that in this work, we do not aim to prove or disprove the existence of biases declared in the experiment by Tomkins et al. (2017). Instead, our focus is on the theoretical validity of the statistical procedures used to conduct such experiments and more generally on principled statistical approach towards designing such experiments.

Finally, we note that the results and tests we discuss in this work are also applicable beyond peer review, and can be used to test for biases in other domains such as admissions and hiring.

The remainder of this paper is organized as follows. In Section 2 we present the problem setting formally and describe the experimental setup of Tomkins et al. (2017). In Section 3 we uncover issues (a) - (e) with their test and setup and illustrate the detrimental effect of such issues through simulations. Next, in Sections 4 and 5 we present a novel non-parametric approach to testing for biases and corresponding statistical tests as well as the alternative design of the experimental procedure. The detailed analysis is given in Section 6. We conclude the paper with a discussion in Section 7.

2 Preliminaries

The general peer-review setup we study for testing biases using single and double blind review is as considered in Tomkins et al. (2017). We study a conference peer-review setup where n papers are submitted at once and m independent reviewers are available to review submissions, where m is assumed to be an even number. With a goal to test whether single-blind reviewing induces a bias against or in favor of some groups of authors, we consider some pre-defined set of k binary mutually non-exclusive properties pertaining to the author(s) of any paper to be tested for bias. For example, a property could be “the first author is female” or “majority of authors are from the USA”. Each paper $j \in [n]$ is then associated with k indicator variables $w_j^{(1)}, \dots, w_j^{(k)}$, where $w_j^{(\ell)} = 1$ if paper j satisfies property ℓ and $w_j^{(\ell)} = -1$ otherwise. For each $\ell \in [k]$ we let $\mathcal{J}_\ell \subseteq [n]$ denote the set of papers that satisfy property ℓ and $\overline{\mathcal{J}}_\ell = [n] \setminus \mathcal{J}_\ell$ denote its complement.¹

¹Here, we adopt the standard notation $[\nu] = \{1, 2, \dots, \nu\}$ for any positive integer ν .

For each property $\ell \in [k]$ we are interested in whether *single-blind peer review setup induces a bias* against or in favor of papers that satisfy this property. For example, if we consider property “the first author is female”, then we aim at testing for the bias against or in favor of papers with female first author. Note that with respect to the properties, the study is observational in that we cannot assign author identities to papers at random. Hence, the effect of confounding is unavoidable and utmost care must be taken to address presence of confounding factors.

For brevity, in the main text we consider the case of a single property of interest ($k = 1$) which captures the complexity of our problem. For ease of notation we drop index ℓ from $w^{(\ell)}$ and \mathcal{J}_ℓ . In Appendix A we generalize the results to $k > 1$. Let us now give details of the testing procedure used by Tomkins et al. (2017).

Experimental setup of Tomkins et al. The peer review process in their experiment is organized as follows. Reviewers are uniformly at random divided into two groups of equal sizes, corresponding to two conditions: (i) Double-Blind condition (DB) in which reviewers do not observe identities of papers’ authors; and (ii) Single-Blind condition (SB) in which reviewers observe identities of the papers’ authors. Next, each paper is assigned to λ reviewers from the SB group and λ reviewers from the DB group such that each reviewer reviews at most μ submissions, where λ and μ are predefined constants. In both conditions, if any reviewer $i \in [m]$ is assigned to any paper $j \in [n]$, then she/he returns a binary accept/reject recommendation and possibly a numeric score that estimates a quality of the paper as perceived by reviewer, accompanied by a textual review.

Model and test used by Tomkins et al. We begin by introducing an idealized version of their model. They assume a parametric, logistic model for the binary decisions made by SB reviewers. Specifically, for each paper $j \in [n]$, let $Y_{1j}, \dots, Y_{\lambda j}$ denote the binary accept/reject decisions given by the λ reviewers assigned to paper j in the SB setup. It is assumed that $\{Y_{rj}\}_{r \in [\lambda]}$ are independent draws from a Bernoulli random variable with an expectation π_j satisfying

$$\log \frac{\pi_j}{1 - \pi_j} = \beta_0 + \beta_1 q_j^* + \beta_2 w_j, \quad (1)$$

where q_j^* is a “true” underlying score of paper j , w_j is an indicator of property satisfaction and $\{\beta_0, \beta_1, \beta_2\}$ are unknown coefficients. In words, the model says that if there is a positive (respectively negative) bias with respect to a property of interest, then the fact that paper satisfies the property increases (respectively decreases) the log-odds of the probability of recommending acceptance by $2\beta_2$ as compared to the case if the same paper does not satisfy the property. The main difficulty with this model in the peer review setting lies in the fact that true scores $\{q_j^*, j \in [n]\}$ are unknown and hence standard tests for logistic regression model are not readily applicable.

In order to overcome the unavailability of true scores $\{q_j^*, j \in [n]\}$ in the model (1), Tomkins et al. (2017) use a plug-in estimate: they replace q_j^* with the mean \tilde{q}_j of scores given by the DB reviewers to paper j , for every $j \in [n]$. Under this approximation and using $\tilde{q}_1, \dots, \tilde{q}_n$, they obtain maximum likelihood estimates of coefficients $\{\hat{\beta}_0, \hat{\beta}_1, \hat{\beta}_2\}$ and then use the standard Wald test (Weisberg, 2005) to test for significance of the coefficient β_2 . A bias is declared present if the coefficient β_2 is found significant; the direction of the bias is determined as the sign of $\hat{\beta}_2$.

3 Problems with the past approach

In this section we identify several issues that should be taken into account when testing for biases in the setup we consider. Noting that the issues themselves are general, we motivate and discuss them in context of the prior work by Tomkins et al. (2017) and investigate possible consequences of these issues through synthetic simulations. In the simulations to follow, we juxtapose algorithm by Tomkins et al. (2017) to our DISAGREEMENT test introduced later in the paper. Complete details of all simulations are given in Appendix D.

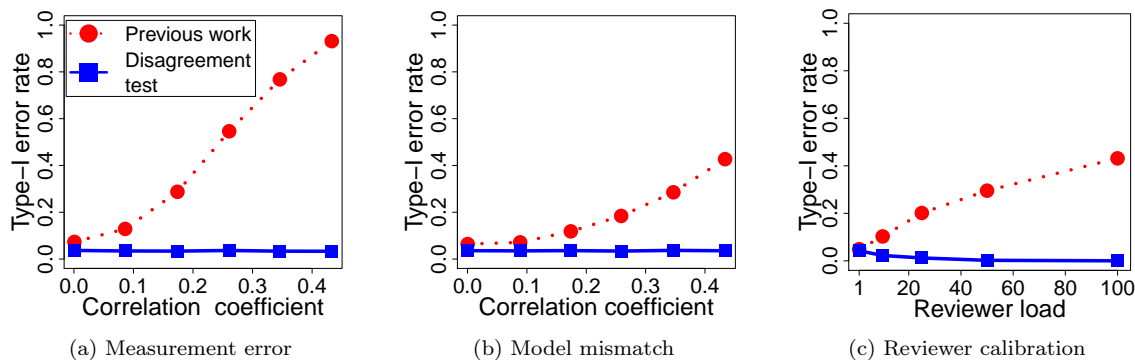


Figure 2: Type-I error of the test from previous work (Tomkins et al. 2017) blows up under three different setups: bias is absent in all simulations and the tests are designed to limit the Type-I error to at most 0.05. In contrast, our DISAGREEMENT test is robust to violations of modelling assumptions. Error bars are too small to be visible.

3.1 Testing procedure

We begin from the issues that are pertinent to the testing procedure used by Tomkins et al. (2017). To this end, recall that with respect to the property of interest the experiment is observational. Hence we *cannot* assume independence between the indicator of property satisfaction w and the true score q^* . Moreover, a non-trivial amount of correlation between some properties is plausible. Consider for example a property “paper has author from top univeristy”. For this property a non-trivial correlation between true scores and indicator of property satisfaction is natural to expect. While correlation itself does not cause issues, we identify three conditions which coupled with correlation can be significantly harmful.

(a) Measurement error. Tomkins et al. (2017) report low interreviewer agreement between DB reviewers which means that the estimates $\hat{q}_1, \dots, \hat{q}_n$ of the true scores by the DB reviewers are noisy. It is known (Stefanski and Carroll, 1985; Brunner and Austin, 2009) that noisy covariate measurement coupled with correlation between some covariates may inflate the Type-I error rate of the Wald test for logistic regression. We now investigate the impact of measurement error on the Type-I error rate of the Tomkins et al. test through simulations. We consider absence of any bias, and assume that model (1) with $\beta_2 = 0$ is correct for both DB and SB reviewers. We consider DB reviewers to report noisy estimates of true scores q_j^* , and vary the correlation between q^* and w . The level of noise was selected to keep correlation between the two DB reviewers assigned to each paper at the level of 0.6, which is much better than the actual interreviewer agreement observed by Tomkins et al. (2017) (correlation 0.37). We plot the Type-I error rates in Figure 2a for the test in Tomkins et al. (2017) and our proposed test, both tests are designed to restrict the Type-I error rate to 0.05.

Figure 2a indicates a strong detrimental effect of measurement error on the validity of the test by Tomkins et al. (2017). Given that interreviewer agreement in the actual WSDM conference experiment was low, the fact that some properties considered by Tomkins et al. may lead to correlations between q^* and w is concerning, because it could potentially undermine the validity of their findings.

The simulations in Section 1 follow the setup presented here: Figures 1a and 1b consider measurement error with correlation fixed at 0.4 (Figure 1a) and 0.6 (Figure 1b) and show that (a) the negative effect of measurement error on the Type-I error rate exacerbates as sample size grows and (b) measurement error may also hinder the power of the test. Figure 1c has zero correlation and no measurement error, satisfying all the assumptions of the test by Tomkins et al.

(b) Model mismatch. Model (1) assumes a specific parametric relationship, which may not hold in practice. In order to check the effect of model mismatches, we consider a violation of the model (1) and suppose that

the correct model for both SB and DB reviewers is

$$\log \frac{\pi_j}{1 - \pi_j} = \beta_0 + \beta_1 (q_j^*)^3 + \beta_2 w_j,$$

that is, instead of expected linear input, true scores of papers appear in the model raised to the power 3. To isolate the effect of model mismatch, we assume that true scores $q_j^*, j \in [n]$, are known exactly to the test of Tomkins et al. and hence abstract out the impact of the measurement error. We again consider an absence of any bias and set $\beta_2 = 0$ for both SB and DB reviewers. We then perform simulations similar to those in item (a). Figure 2b shows the results of the simulations.

(c) Reviewer calibration. The test employed by Tomkins et al. (2017) treats reviews given by the same reviewer as independent. In practice this assumption may be violated due to correlations introduced by reviewer’s calibration (Wang and Shah, 2018). While some easy calibrations such as harshness/leniency can be captured by simple parametric extensions of model (1), more subtle patterns are beyond the scope of this model. Suppose for example that the strength of reviewers’ input depends on paper’s clarity — the better the paper is written, the lower the contribution due to reviewers’ calibration. Assume also that we are given a set of papers such that true score of each paper is proportional to the clarity of the paper (we formalize construction in Appendix D.1.3). Coupled with the correlation between q^* and w , this pattern is sufficient to break Type-I error guarantees of the test of Tomkins et al. Again, to isolate the impact of reviewers’ calibration, we assume that (i) true scores $q_j^*, j \in [n]$, are known to the test by Tomkins et al. and (ii) model (1) is marginally correct for each reviewer, that is, each reviewer follows model (1) for each paper she/he reviews, but her/his decisions for different papers are correlated in a specific way.

Figure 2c shows a result of simulations in which we vary the number of papers per reviewer, keeping correlation between q^* and w fixed at 0.75 and the total number of papers fixed at $n = 1000$. We simulate a wide range of reviewer load μ including small to medium loads of 5-15 papers typical in machine learning conferences like NeurIPS and larger loads of 40 or higher found in other smaller conferences.

3.2 Experimental setup

The issues discussed above pertain to the testing procedure and modelling assumptions made by Tomkins et al. (2017). We now issue a commentary regarding the experimental setup considered in their work which comprises a random partition of reviewers into SB and DB groups within a standard peer review procedure. In particular, we show that the setup itself may create problems in controlling the Type-I error.

(d) Non-blind bidding. In the experiment by Tomkins et al. papers are allocated to reviewers based on preferences (“bids”) declared by reviewers (reviewers could indicate that they want to review some papers and do not want to review others). Importantly, the reviewers in the SB setup also get to see author identities in the bidding stage, which may act as a confounding factor in tests for bias in the acceptance/rejection of papers. This is indeed pointed out as a caveat by Tomkins et al. (2017) in their paper.

To illustrate the possible effect, consider a property of interest “paper has a famous author” and suppose that among all reviewers there is a subset of lenient reviewers who additionally want to read papers from top authors with the hope of reading better papers. Then in DB setup such reviewers cannot use author identity information and hence make their bidding decisions based on title and abstract only; in contrast, in SB setup these reviewers tend to bid on papers authored by top authors. Given that reviewers who preferentially bid on papers with top authors in SB condition are by coincidence lenient, the difference in bidding behavior may result in structurally different evaluations between conditions even when reviewers’ evaluations are unbiased, leading to a blow-up of the Type-I error rate of any reasonable test. Figure 3a shows a result of simulations (formal setup is in Appendix D.1.4) in which we compare non-blind and blind bidding conditions for SB reviewers and indicates a possible detrimental effect of non-blind bidding.

(e) Reviewer assignment. One might imagine that a natural requirement to conduct the bidding in a double blind fashion for both DB and SB reviewers would fix the issues with the setup of Tomkins et al. However, perhaps surprisingly, we show that even if both groups bid in a double blind fashion (or even if the bidding process is eliminated entirely), and even if the reviewers are assigned to DB or SB groups

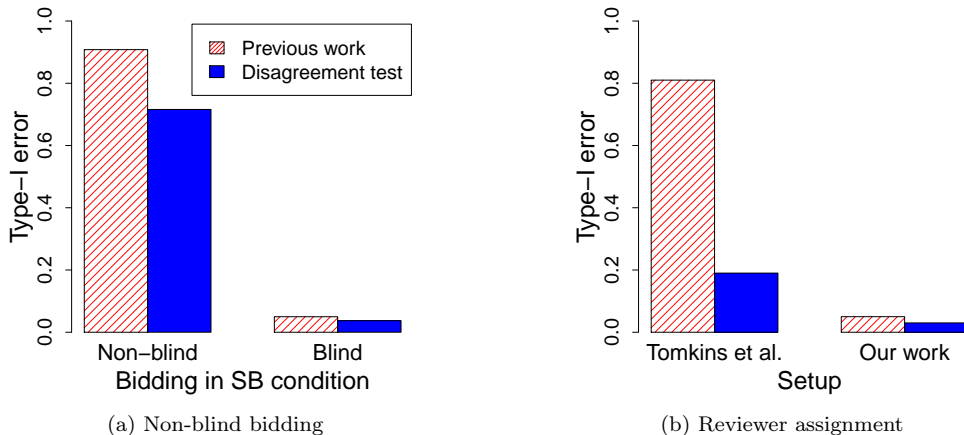


Figure 3: The experimental setup from previous work (Tomkins et al. 2017) violates Type-I error guarantees of testing procedures. Bias is absent in all simulations and the tests are designed to limit the Type-I error to at most 0.05. Note that the issues which pertain to the experimental setup rather than the modelling break guarantees of both tests (leftmost columns). In contrast, our proposed setup with fully blind bidding procedure and careful management of assignment ensures Type-I error guarantees for both tests (rightmost columns). Error bars are too small to be visible.

uniformly at random, the non-random assignment using algorithms such as TPMS (Charlin and Zemel, 2013) that assigns reviewers to papers maximizing some notion of “similarities” can still lead to a violation of the Type-I error guarantees. We give a formal construction in Appendix D.1.5; the intuition is as follows. Quoting Lamont (2009), “evaluators often define excellence as «what speaks to me» which is akin to «what is most like me»”, that is, a similarity between a paper and a reviewer may influence the decision. That said, we construct these similarities in a careful manner: our choice ensures that despite reviewers being allocated to DB or SB conditions at random, the popular TPMS assignment algorithm with high probability constructs assignments that are in some sense structurally different between SB and DB conditions, which in turn leads to structurally different evaluations. Our construction, along with correlation between q^* and w , introduces spurious correlations in the data, thereby violating some key independence assumptions and leading to the inflation of Type-I error.

Figure 3b shows a result of simulations in which we compare the setup of Tomkins et al. (2017) with our proposed experimental setup introduced in Section 5.2. Notably, under the setup of Tomkins et al. even the DISAGREEMENT test which is robust to various issues discussed in Section 3.1 is unable to control for the Type-I error. In contrast, observe that under our proposed experimental setup, both the DISAGREEMENT test and the test by Tomkins et al. control for the Type-I error rate at the desired level.

Importantly, we underscore that while our experimental procedure mitigates the issues with the experimental setup of Tomkins et al., their test is still *susceptible to the issues we discussed in Section 3.1* even under our experimental setup. Finally, under the setup of Tomkins et al. the phenomenon of the Type-I guarantee violation is not restricted to the TPMS assignment and can occur in a much broader class of reviewer assignment algorithms.

4 Novel framework to test for biases

In Section 3 we identified five key limitations of the approach taken by Tomkins et al. (2017). Three of these limitations pertain to the testing procedure and the two limitations relate to the design of the experiment itself. In the next sections we design a set of tests and experimental setup with strong guarantees, and which

overcome the aforementioned limitations. In this section we begin from principled definition of a bias testing problem that generalizes one made by Tomkins et al. and does not make any restrictive assumptions.

At a high level, our approach to testing for biases is different from those proposed by Tomkins et al. in two ways. First, we relax two strict modelling assumptions: (i) instead of assuming existence of true qualities of submissions, we allow subjectivity in reviewer evaluations (Kerr et al., 1977; Ernst and Resch, 1994; Bakanic et al., 1987; Mahoney, 1977; Lamont, 2009; Noothigattu et al., 2018), and (ii) we do not assume any specific form of the relationship between a paper and its probability of acceptance by a reviewer. Instead, we allow these probabilities to be completely arbitrary and define the bias in terms of these probabilities. Second, we treat this problem conceptually differently from the work of Tomkins et al. The test therein treats the problem as that of one-sample testing and uses DB scores as a plugin estimate of true scores in SB model. In contrast, we approach this problem through the lenses of two-sample testing, where SB and DB reviews form the two samples, and the goal is to test whether they belong to the same distribution. This perspective helps us to avoid a number of issues discussed in Section 3.

Formally, let $\Pi^{\text{db}} \in [0, 1]^{m \times n}$ be a matrix whose $(i, j)^{\text{th}}$ entry, denoted as $\pi_{ij}^{(\text{db})}$, represents a probability that reviewer i would recommend acceptance of paper j if that paper is assigned to that reviewer in DB setup. Similarly, let matrix $\Pi^{\text{sb}} \in [0, 1]^{m \times n}$ be an analogous matrix in SB setup, and denote its $(i, j)^{\text{th}}$ entry as $\pi_{ij}^{(\text{sb})}$.

Let \mathcal{R}_{SB} be the set of reviewers allocated to the SB condition. Moreover, for each $i \in \mathcal{R}_{\text{SB}}$, let $\mathcal{P}_{\text{SB}}(i)$ denote the set of papers assigned to reviewer i and let $Y_{ij} \in \{0, 1\}$ denote the accept/reject decision given by reviewer i for paper $j \in \mathcal{P}_{\text{SB}}(i)$. We similarly define set of DB reviewers \mathcal{R}_{DB} and their decisions $\{X_{ij} : i \in \mathcal{R}_{\text{DB}}, j \in \mathcal{P}_{\text{DB}}(i)\}$. We are interested in testing for biases with respect to a property of interest. To this end, recall our notation $\mathcal{J} \subseteq [n]$ for the set of papers that satisfy a property of interest, and $\bar{\mathcal{J}}$ as its complement.

With this notation in place, we now define two formulations of the bias testing problem — “absolute” and “relative”: The relative bias setting is strictly more general than the absolute bias setting, but also leads to more restrictive results. Importantly, the tests we will introduce in Section 5.1 are applicable to both formulations without additional modifications.

4.1 Absolute bias problem

In the absence of bias, the knowledge of authors’ identities does not induce any difference in reviewers’ behaviour. In the biased hypothesis, there is a positive bias in favor of papers that satisfy a property of interest: reviewers in SB condition are more lenient towards papers from \mathcal{J} and more harsh towards papers from $\bar{\mathcal{J}}$ than they would be in DB condition. The following problem formalizes this intuition.

Problem 1 (Absolute bias problem). Given significance level $\alpha \in (0, 1)$ and decisions of SB and DB reviewers, the goal is to test the following hypotheses:

$$\begin{aligned} H_0 &: \forall i \in [m] \forall j \in [n] \quad \pi_{ij}^{(\text{sb})} = \pi_{ij}^{(\text{db})} \\ H_1 &: \forall i \in [m] \forall j \in [n] \quad \begin{cases} \pi_{ij}^{(\text{sb})} \geq \pi_{ij}^{(\text{db})} & \text{if } j \in \mathcal{J} \\ \pi_{ij}^{(\text{sb})} \leq \pi_{ij}^{(\text{db})} & \text{if } j \in \bar{\mathcal{J}}, \end{cases} \end{aligned} \quad (2)$$

where at least one inequality in the alternative hypothesis (2) is strict.

Note that one can define an alternative that represents a bias against papers from \mathcal{J} simply by exchanging the sets \mathcal{J} and $\bar{\mathcal{J}}$ in (2). Our goal is to design a testing procedure that controls for Type-I error and has non-trivial power for any pair of matrices $\Pi^{\text{sb}}, \Pi^{\text{db}}$ that fall under definition of Problem 1.

Non-trivial power. Informally, we say that the test has non-trivial power if for choices of Π^{sb} and Π^{db} for which the presence of bias is “obvious”, the test is able to detect the bias with probability that goes to 1 as number of papers in both \mathcal{J} and $\bar{\mathcal{J}}$ grows to infinity. Formally, we say that matrices Π^{sb} and Π^{db} satisfy alternative hypothesis (2) with margin δ , if all inequalities in equation (2) are satisfied with margin $\delta > 0$,

that is, $|\pi_{ij}^{(\text{sb})} - \pi_{ij}^{(\text{db})}| > \delta \quad \forall (i, j) \in [m] \times [n]$. Then we say that the testing procedure has non-trivial power if for any $\varepsilon > 0$ and for any $\delta > 0$ there exists $n_0 = n_0(\varepsilon, \delta)$ such that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then for any Π^{sb} and Π^{db} that satisfy alternative hypothesis (2) with margin δ , the power of testing procedure is at least $1 - \varepsilon$.

For instance, if the logistic model (1) is correct for both SB and DB reviewers for some $\beta_0^{(\text{sb})} = \beta_0^{(\text{db})} = \beta_0$, $\beta_1^{(\text{sb})} = \beta_1^{(\text{db})} = \beta_1 > 0$, $\beta_2^{(\text{db})} = 0$ and $|\beta_2^{(\text{sb})}| > 0$, then the requirement of non-trivial power ensures that for any choice of true scores bounded in absolute value by a universal constant and any choice of property satisfaction indicators, the test has power growing to 1 as $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ goes to infinity.

4.2 Relative bias problem

In Problem 1 we assumed that SB (or DB) condition itself does not cause any change in reviewers' behaviour. We now consider a generalization of Problem 1 which accommodates an additional confounding factor — a bias in the reviewer simply due to her/his assignment in the SB or the DB group (and independent of the paper or its characteristics). For example, reviewers may not have any bias with respect to the property of interest, but just being placed in the SB condition may induce more harsh opinions than the reviewers in DB. Formally, recall the null hypothesis $\pi_{ij}^{(\text{sb})} = \pi_{ij}^{(\text{db})} \quad \forall (i, j) \in [m] \times [n]$ in Problem 1. Instead, under the null, we now allow $\pi_{ij}^{(\text{sb})} = f_0(\pi_{ij}^{(\text{db})})$, for some non-decreasing function $f_0 : [0, 1] \rightarrow [0, 1]$. Of course, one may not know the function f_0 and the goal of this general problem is to design a test that is guaranteed to control over Type-I error and has non-trivial power uniformly for all functions f_0 that belong to some set of non-decreasing functions \mathcal{F} .

Problem 2 (Relative bias problem). Given significance level $\alpha \in (0, 1)$, class of functions \mathcal{F} and decisions of SB and DB reviewers, the goal is to test the following hypotheses:

$$\begin{aligned} H_0 : \forall i \in [m] \forall j \in [n] \quad \pi_{ij}^{(\text{sb})} &= f_0(\pi_{ij}^{(\text{db})}) \\ H_1 : \forall i \in [m] \forall j \in [n] \quad &\begin{cases} \pi_{ij}^{(\text{sb})} \geq f_0(\pi_{ij}^{(\text{db})}) & \text{if } j \in \mathcal{J} \\ \pi_{ij}^{(\text{sb})} \leq f_0(\pi_{ij}^{(\text{db})}) & \text{if } j \notin \mathcal{J} \end{cases}, \end{aligned} \quad (3)$$

where f_0 is some unknown function from \mathcal{F} and at least one inequality in the alternative hypothesis (3) is strict.

For example, if the logistic model (1) is correct for both SB and DB reviewers (with $\beta_2^{(\text{db})} = 0$), but intercepts β_0 in SB and in DB conditions are allowed to be different, then the corresponding matrices Π^{sb} and Π^{db} do not fall under the definition of Problem 1, but can be captured by Problem 2 with specific choice of \mathcal{F} as we will discuss in Section 6.2.

The definition of non-trivial power transfers to the relative bias problem with the exception that all $\pi_{ij}^{(\text{db})}$ are substituted by $f_0(\pi_{ij}^{(\text{db})})$ for $f_0 \in \mathcal{F}$. Our goal is to design a testing procedure that controls for Type-I error and has non-trivial power for any pair of matrices $\Pi^{\text{sb}}, \Pi^{\text{db}}$ that fall under definition of Problem 2 for any function $f_0 \in \mathcal{F}$. Ideally, we would like to achieve this goal for a set of functions \mathcal{F} that contains all non-decreasing functions $f : [0, 1] \rightarrow [0, 1]$.

5 Proposed solution

We now introduce the proposed experimental setup as well as statistical tests we study in this work. We subsequently analyze them in the context of Problems 1 and 2 in Section 6.

5.1 Testing procedures

In order to avoid correlations introduced by reviews given by the same reviewer, our tests use at most one decision per reviewer. As we discuss in Section 5.2, we do so by first matching reviewers into pairs, consisting

of one SB and one DB reviewer who review a common paper. For the moment, assume that we are given a set of tuples \mathcal{T} , where each tuple $t \in \mathcal{T}$ consists of a paper $j_t \in [n]$, decision of a SB reviewer for this paper Y_{j_t} , decision of a DB reviewer for this paper X_{j_t} and indicator of property satisfaction w_{j_t} , with a constraint that each reviewer contributes her/his decision to at most one tuple. With this notation, we now present two tests we consider in this work. As we show subsequently, either of these tests would suffice for the absolute bias problem, but for the relative bias problem they cater to different models of reviewers' behaviour with non-intersecting areas of applicability. To provide intuition behind the tests, we define them in context of the absolute bias problem (Problem 1) and discuss their applicability to the relative bias problem later.

Disagreement-based test. A high-level idea of the test is as follows. Consider a pair of SB and DB reviewers who disagree in their decisions for some paper. Then under the null hypothesis, the events “SB accepts and DB rejects” and “SB rejects and DB accepts” are equally likely. In contrast, if the null hypothesis is violated, then depending on the property satisfaction and the direction of the bias, SB reviewer is more (or less) likely to vote for acceptance than her/his DB counterpart.

We now formally present the DISAGREEMENT test as Test 1. In Step 1 two empty arrays U and V are initialized. Next, in Step 2 we focus on pairs of SB and DB reviewers disagreeing in their decisions for a paper they both review. For each of the corresponding tuples, we add the decision of SB reviewer to the array U if a paper satisfies the property of interest and to V otherwise. Finally, in Step 3 we define a test statistic τ . According to the aforementioned intuition, under the null hypothesis τ should be close to 0, but under the alternative it should be large in absolute value. Hence, to make a decision we run a permutation test and reject the null in Step 3 if this test suggests that $|\tau|$ is too large for a given significance level α .

Test 1 DISAGREEMENT

Input: Significance level $\alpha \in (0, 1)$

Set of tuples \mathcal{T} , where each $t \in \mathcal{T}$ is of the form $(j_t, Y_{j_t}, X_{j_t}, w_{j_t})$ for some paper $j \in [n]$.

1. Initialize U and V to be empty arrays.

2. For each tuple $t \in \mathcal{T}$, if $Y_{j_t} \neq X_{j_t}$, append Y_{j_t} to $\begin{cases} U & \text{if } w_{j_t} = 1 \\ V & \text{if } w_{j_t} = -1 \end{cases}$.

3. Run a permutation test at the level α to test if entries of U and V are exchangeable random variables, using the test statistic:

$$\tau = \frac{1}{|U|} \sum_{r \in [|U|]} U_r - \frac{1}{|V|} \sum_{r \in [|V|]} V_r.$$

4. Reject the null if and only if the permutation test rejects the null. (If either of the arrays V and U is empty, the test keeps the null.)

Counting-based test. The test is built on a simple intuition. Assume for the moment that SB setup induces a bias against papers from \mathcal{J} and a bias in favor of papers from $\overline{\mathcal{J}}$. Then it is likely that papers from \mathcal{J} will receive less number of positive recommendations in SB setup as compared to DB setup. Symmetrically, for papers from $\overline{\mathcal{J}}$ we expect reviewers in SB to be more lenient than their DB counterparts. In contrast, if there is no bias at all, then we expect the aforementioned differences to be small.

We now formally present the COUNTING test as Test 2. In Step 1 two empty arrays U and V are created which in Step 2 are populated with differences between decisions of SB and DB reviewers for papers from \mathcal{J} and $\overline{\mathcal{J}}$ respectively. Importantly, in contrast to the DISAGREEMENT test, in the COUNTING test we do not condition on disagreeing pairs of reviewers. Noticing that mean value of entries of U (respectively V) measures the change of attitude towards papers from \mathcal{J} (respectively $\overline{\mathcal{J}}$) between SB and DB conditions, in Step 3 we compute a test statistic γ which compares these changes. According to the aforementioned intuition, under correct null hypothesis the test statistic should be close to 0. Finally, in Step 4 we make a decision using concentration properties of the test statistic.

Test 2 COUNTING

Input: Significance level $\alpha \in (0, 1)$ Set of tuples \mathcal{T} , where each $t \in \mathcal{T}$ is of the form $(j_t, Y_{j_t}, X_{j_t}, w_{j_t})$ for some paper $j \in [n]$.1. Initialize U and V to be empty arrays.2. For each tuple $t \in \mathcal{T}$, append $(Y_{j_t} - X_{j_t})$ to $\begin{cases} U & \text{if } w_{j_t} = 1 \\ V & \text{if } w_{j_t} = -1 \end{cases}$.3. If either of the arrays V and U is empty, keep the null and terminate. Otherwise, set the test statistic γ as follows:

$$\gamma = \frac{1}{|U|} \sum_{r \in [|U|]} U_r - \frac{1}{|V|} \sum_{r \in [|V|]} V_r.$$

4. Reject the null hypothesis if and only if

$$|\gamma| > \sqrt{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha}.$$

5.2 Setup of the experiment

We now propose the setup of the experiment to overcome the issues highlighted in Section 3.2 and discuss a construction of the set \mathcal{T} used by the tests introduced above. At a higher level, the proposed setup has two main differences from one considered by Tomkins et al. (2017). First, bidding is performed in blind manner by both SB and DB reviewers (Step 1 below). Second and more importantly, to avoid issues caused by non-random reviewers' assignment, we perform paper assignment and reviewer allocation to conditions jointly in a carefully selected manner (Steps 2-4 below).

We now formally present the experimental procedure as Procedure 1. It takes as input parameters of paper and reviewer loads together with any assignment algorithm that operates on similarities and/or bids. In Step 1 reviewers bid on the papers in blind manner, that is, using only title and abstract of submissions. Notice that in contrast to the Tomkins et al. setup, bidding happens even before the reviewers are allocated to SB or DB conditions. In Step 2 we find a partial assignment of papers to reviewers which satisfies $(\lambda = 2, \mu = 1)$ -load constraints. Depending on the relationship between n and m , we may include only a subset of papers or reviewers in this assignment. For example, in case 2b we do not have enough reviewers to respect the one paper per reviewer constraint and hence we select subset of papers of appropriate size such that it includes approximately equal number of papers from \mathcal{J} and $\bar{\mathcal{J}}$ and find the assignment for selected papers only. The constraint on the number of papers from \mathcal{J} and $\bar{\mathcal{J}}$ is to ensure that the resulting set \mathcal{T} is balanced which is necessary for non-trivial power. The corresponding assignment A^* is a building block for our tests which will use the reviews from this assignment only. Next, in Step 3 reviewers are allocated to conditions in a specific manner which is crucial for our statistical guarantees. In Step 4 we find a full assignment A that is a completion of the partial assignment A^* , meaning that if reviewer i was assigned to paper j in assignment A^* , she/he is also assigned to this paper in A . Finally, in Step 5 we construct a set of tuples \mathcal{T} that is used by the DISAGREEMENT and COUNTING algorithms in Step 6. Importantly, by construction we ensure that each reviewer contributes at most one decision to the set \mathcal{T} .

As we show below, the experimental Procedure 1 overcomes the issues with the experimental setup we discussed in Section 3.2 and leads to provable control over Type-I error for our DISAGREEMENT and COUNTING algorithms. We underscore that (i) DISAGREEMENT and COUNTING tests are not tied to particular experimental procedure we introduce and can be applied under the setup of Tomkins et al. (2017) with caveats discussed in Section 3.2. For instance, in simulations (a)-(d) of Section 3 the DISAGREEMENT test was applied under the setup of Tomkins et al. More details on this remark are provided in Appendix B; (ii) as requested by the DISAGREEMENT and COUNTING tests, the set of tuples \mathcal{T} constructed in Step 5 contains

at most one decision of each reviewer. This requirement allows our tests to be agnostic to reviewer calibration which may otherwise undermine Type-I error guarantees as demonstrated in Section 3.1. However, if one treats reviews given by the same reviewer as independent, thereby ignoring issues with reviewer calibration, then in Step 5 of Procedure 1 one can construct a larger set \mathcal{T} by using full assignment A and allowing each reviewer to contribute multiple decisions to the set \mathcal{T} .

Procedure 1 Design of the experiment

Input: Paper load $\lambda \geq 1$
Reviewer load $\mu \geq 1$
Assignment algorithm \mathcal{A}

1. Reviewers bid on papers in blind manner
2. Depending on the relationship between number of papers (n) and reviewers (m):
 - (a) If $m > 2n$, select $2n$ reviewers uniformly at random and use algorithm \mathcal{A} to assign each paper to 2 reviewers from the selected pool such that each reviewer is assigned to one paper
 - (b) If $m < 2n$, select $m/2$ papers such that proportions of papers from \mathcal{J} and $\overline{\mathcal{J}}$ are as close to each other as possible. Use algorithm \mathcal{A} to assign each selected paper to 2 reviewers such that each reviewer is assigned to one paper
 - (c) If $m = 2n$, use algorithms \mathcal{A} to assign each paper to 2 reviewers such that each reviewer is assigned to one paper

Denote the corresponding assignment as A^*

3. For each paper in assignment A^* , allocate one assigned reviewer to DB condition and another assigned reviewer to SB condition uniformly at random. If at this point there are reviewers who are not allocated to conditions, allocate half of them to SB and half to DB uniformly at random
 4. Using algorithm \mathcal{A} , complement assignment A^* such that each paper is assigned to λ SB and λ DB reviewers and each reviewer reviews at most μ papers. Denote the corresponding assignment as A and begin review process according to this assignment
 5. When the review process is finished, construct a set \mathcal{T} as follows. For every paper j from the assignment A^* and corresponding pair (i_1, i_2) of SB and DB reviewer, add tuple $(j, Y_{i_1 j}, X_{i_2 j}, w_j)$ to the set \mathcal{T}
 6. Run statistical test on the set \mathcal{T}
-

Finally, in addition to facilitating the experiment, in the interest of fairness in the review (Tomkins et al., 2017) the experimental procedure should ensure that in the eventual assignment each paper is reviewed by equal number of SB and DB reviewers. By construction, Procedure 1 satisfies this requirement: Step 4 ensures that in the final assignment A each paper is assigned to λ SB reviewers and λ DB reviewers.

6 Analysis

We now present the analysis of the COUNTING and DISAGREEMENT tests in context of absolute and relative bias problems.

6.1 Absolute bias problem

We begin our analysis from the absolute bias problem and first formulate the main theorem of this section.

Theorem 1. *For any significance level $\alpha \in (0, 1)$, under the setup of the absolute bias problem (Problem 1), let the experiment be organized according to Procedure 1. Then the DISAGREEMENT and COUNTING tests are guaranteed to control for Type-I error at the level α , and also satisfy the requirement of non-trivial power.*

Remark. 1. As demonstrated in Figure 4, In practice, the DISAGREEMENT test has a higher power as compared to the COUNTING test and should be employed under conditions of the absolute bias problem.

2. Notice that the outcomes of the DISAGREEMENT and COUNTING tests depend on a set \mathcal{T} provided to the tests as input. That is, for two different sets \mathcal{T}_1 and \mathcal{T}_2 (that for example correspond to different assignments A_1^* and A_2^* constructed in Step 2 of Procedure 1), the outcomes of the tests might be different. Hence, one should fix a set \mathcal{T} before observing reviewers’ decisions to avoid chasing statistical significance.

3. Finally, the DISAGREEMENT and COUNTING tests are also applicable to the experimental procedure used by Tomkins et al. (2017) and are guaranteed to be robust to issues (a)-(c) from Section 3.1. The formal statement is given in Appendix B.

We now discuss the issues (a)-(e) considered in Section 3 in the context of our DISAGREEMENT and COUNTING tests.

- **Noise.** The DISAGREEMENT and COUNTING tests do not rely on any estimation of papers’ qualities made by reviewers. Moreover, we do not even assume that there exists some objective quantity that can be estimated. Hence, our tests do not suffer from issues caused by noisy estimates of scores given by DB reviewers as illustrated by Figure 2a in case of the DISAGREEMENT test.
- **Model mismatch.** The only assumption we make is that under correct null hypothesis there is no difference in behavior of SB and DB reviewers. Hence, Theorem 1 guarantees that our tests are robust to violations of specific parametric model (1) as illustrated by Figure 2b.
- **Reviewer calibration.** We circumvent the detrimental effect of correlations introduced by reviewers’ calibration by requiring that each reviewer contributes at most one review to the test. See Figure 2c for an illustration. Of course, such robustness comes at the cost of some power, but we notice that our matching procedures guarantee the use of at least a constant fraction of available data, thereby limiting reduction in the power.
- **Non-blind bidding.** The issue with bidding is straightforwardly resolved by requesting blind bidding from both SB and DB reviewers. As illustrated by Figure 3a, we abstract out possible confoundings due to difference in bidding behaviour and ensure that the observed difference in decisions (if any) is due to bias in evaluations and not in bidding.
- **Reviewer assignment.** The experimental design proposed in Procedure 1 allows to execute any assignment algorithm without breaking the guarantees of the tests as demonstrated in case of TPMS assignment by Figure 3b. The key part of the procedure that ensures such robustness are Steps 2 and 3 where we first find triples of two reviewers and a paper they review and then randomly allocate one reviewer from each triple to SB and one to DB. In this manner, we ensure that parts of the final assignment A that are used for testing do not exhibit any structural difference caused by non-random assignment.

We conclude the section with brief discussion of the power of the tests we introduced in this section. To provide fair comparison of tests, the simulations are performed under the setup of the experiment by Tomkins et al. (2017) and the input to the DISAGREEMENT and COUNTING tests is computed according to procedure described in Appendix B. Figure 4 contrasts the powers of the tests under the logistic model (1) in two cases. In Figure 4a, DB reviewers estimate the true scores of the papers with some noise, which in presence of correlation dramatically decreases the power of the test by Tomkins et al. As discussed above, the DISAGREEMENT and COUNTING tests are robust to issues caused by measurement error, as seen in Figure 4a.

In contrast, Figure 4b considers the case when DB reviewers estimate true scores without noise, that is, true scores are known to the Tomkins et al. test. Although in this case their test has the highest power among 3 tests under consideration, we notice that the margin between their test and the DISAGREEMENT test is not as large as in Figure 4a. Notice also that the test by Tomkins et al. gains power by overfitting to the strict model (1) which leads to higher power when the model is correct, but at the cost of not being able to control over Type-I error rate under reasonable violations of the modelling assumptions as discussed in Section 3.

Finally, notice that the COUNTING test relies on the sub-Gaussian approximation to define a threshold and consequently has lower power than the DISAGREEMENT test in both cases. While this suggests that

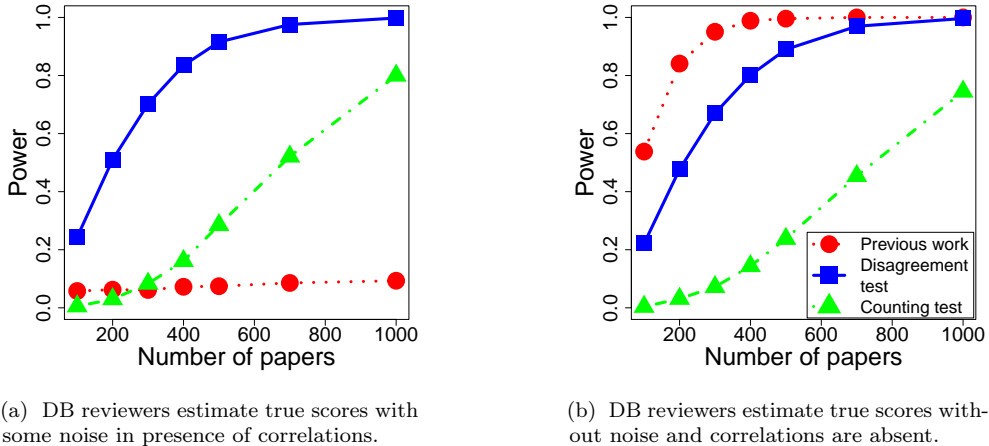


Figure 4: Synthetic simulations evaluating performance of the test in Tomkins et al. (2017) (“previous work”) and the tests introduced in this section (“DISAGREEMENT and COUNTING tests”) in presence of bias when logistic model (1) of reviewers is correct. Larger values are better. Details of the simulation setup are provided in Appendix D.2. Error bars are too small to be visible.

for absolute bias problem the DISAGREEMENT test dominates the COUNTING test, we will show in the next section that under the relative bias problem these tests are incomparable.

6.2 Relative bias problem

In Section 6.1 we showed that if under the absence of bias the behaviour of reviewers in SB and DB conditions is the same, then the DISAGREEMENT and COUNTING tests control for Type-I error and have non-trivial power thus leading to a reliable testing procedure. In this section we relax that assumption and consider a relative bias problem with two specific choices of \mathcal{F} that correspond to popular linear and logistic models. We first show that for each of these choices either the DISAGREEMENT or the COUNTING test leads to a reliable testing. We then conclude the analysis with a negative result saying that no test can control for Type-I error and have non-trivial power over both choices of \mathcal{F} .

Let us now introduce these two choices of \mathcal{F} that we consider throughout the remaining part of this section. To this end, for each $j \in [n]$ we let $q_j \in \mathbb{R}$ denote an unknown “representation” of the paper j , where by representation we imply any function of a paper’s content that defines reviewers’ perception of a paper. For example, it could be that $q_j = q_j^*$, where q_j^* is a true score as defined by Tomkins et al.

Generalized linear model. Under the generalized linear model, the SB condition itself induces a change in reviewers’ evaluations, making them more harsh (or lenient). Moreover, under absence of bias the change in behaviour between SB and DB conditions is described by a constant shift in probability of acceptance for all papers, irrespective of whether they satisfy a property of interest or not. Formally, consider a fixed constant $\Delta \in (0, 0.5)$. The generalized linear model assumes that (i) for every $j \in [n]$ a corresponding representation q_j belongs to the interval $(\Delta, 1 - \Delta)$ and (ii) under absence of bias for every $(i, j) \in [m] \times [n]$ the behavior of reviewer i if she/he reviews paper j is described by the following parametric equations:

$$\text{DB: } \pi_{ij}^{(\text{db})} = q_j \tag{4a}$$

$$\text{SB: } \pi_{ij}^{(\text{sb})} = \nu + q_j, \tag{4b}$$

for some unknown constant $\nu \in (-\Delta, \Delta)$. Provided that matrix Π^{db} was generated according to the model (4a) of DB reviewers, the generalized linear model corresponds to an instance of a relative bias problem with a set

of functions \mathcal{F}_Δ associated to a fixed constant Δ and defined as

$$\mathcal{F}_\Delta = \left\{ h_\nu(t) : (\Delta, 1 - \Delta) \rightarrow [0, 1] \mid \nu \in (-\Delta, \Delta) \right\}, \quad (5)$$

where

$$h_\nu(t) = t + \nu, \quad t \in (\Delta, 1 - \Delta). \quad (6)$$

Indeed, observe that if for any $(i, j) \in [n] \times [m]$ the probability of acceptance $\pi_{ij}^{(\text{db})}$ was generated according to the model (4a), then $\pi_{ij}^{(\text{sb})}$ defined as $\pi_{ij}^{(\text{sb})} = h_\nu(\pi_{ij}^{(\text{db})})$ satisfies model (4b).

Generalized logistic model. Similar to the generalized linear model, under the generalized logistic model the SB condition also induces a change in reviewers' evaluations, but the change now is described as a constant shift in space of log-odds of the acceptance probabilities. Formally, consider a fixed constant $\tilde{\Delta} > 0$. The generalized logistic model assumes that (i) for every $j \in [n]$ a corresponding representation q_j belongs to the interval $(-\tilde{\Delta}, \tilde{\Delta})$ and (ii) under absence of bias for every $(i, j) \in [m] \times [n]$ behaviour of reviewer i if she/he reviews paper j is described by the following parametric equations:

$$\text{DB: } \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} = \beta_0 + \beta_1 q_j \quad (7a)$$

$$\text{SB: } \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = \beta_0 + \tilde{\nu} + \beta_1 q_j, \quad (7b)$$

for some unknown constant $\tilde{\nu} \in (-\tilde{\Delta}, \tilde{\Delta})$, where unknown coefficients β_0 and β_1 are also bounded in absolute value by $\tilde{\Delta}$ and $\beta_1 > 0$. Provided that matrix Π^{db} is generated according to the model (7a) of DB reviewers, one can verify that the generalized logistic model corresponds to an instance of a relative bias problem with a set of functions $\tilde{\mathcal{F}}_{\tilde{\Delta}}$ associated to a fixed constant $\tilde{\Delta}$ and defined as

$$\tilde{\mathcal{F}}_{\tilde{\Delta}} = \left\{ g_{\tilde{\nu}}(t) : [0, 1] \rightarrow [0, 1] \mid \tilde{\nu} \in (-\tilde{\Delta}, \tilde{\Delta}) \right\}, \quad (8)$$

where

$$g_{\tilde{\nu}}(t) = \frac{te^{\tilde{\nu}}}{1 - t + te^{\tilde{\nu}}}, \quad t \in [0, 1]. \quad (9)$$

Indeed, observe that if for some $(i, j) \in [n] \times [m]$ the probability that reviewer i accepts paper j in DB setup $\pi_{ij}^{(\text{db})}$ is generated according to the model (7a), then setting $\pi_{ij}^{(\text{sb})} = g_{\tilde{\nu}}(\pi_{ij}^{(\text{db})})$ we ensure that $\pi_{ij}^{(\text{sb})}$ satisfies model (7b).

The models we defined follow an objective parametric approach assumed by Tomkins et al. (2017) with two differences: (i) our models do not assume that q_j has a known meaning or that it can be measured (for instance, it may be that $q_j = q_j^*$, or that $q_j = (q_j^*)^3$, or q_j may be a complex function of the content of the paper) and (ii) our models do not assume that the bias is described by a linear shift in space of probabilities or log-odds, and instead consider non-parametric definition of the bias as specified in the alternative hypothesis (3).

6.2.1 Positive results

We now show that the COUNTING and DISAGREEMENT tests lead to a reliable testing under the generalized linear and generalized logistic models respectively. Before we formulate the main result of this section, let us provide some intuition behind the models and the corresponding tests.

Generalized linear model. A natural strategy to test for biases under the generalized linear model is to estimate the shift in reviewers' behaviour on papers that belong to \mathcal{J} and to $\bar{\mathcal{J}}$ separately and then compare

the estimates. In fact, the COUNTING testing procedure introduced above as Test 2 follows this strategy and, as guaranteed by Theorem 2, leads to a reliable testing under the generalized linear model.

Generalized logistic model. Intuitively, estimating a constant shift in reviewers’ behavior as done by the COUNTING test may not be the optimal strategy under the generalized logistic model, as under absence of bias the change of behaviour between SB and DB conditions is given by a constant shift in *log-odds* space and not in *probability* space. However, it turns out that the DISAGREEMENT test is able to capture the constant shift in *log-odds* space and hence we still can perform reliable testing under this model, using the DISAGREEMENT algorithm.

Theorem 2. *For any significance level $\alpha \in (0, 1)$, let the experiment be organized according to Procedure 1. Then*

- (a) *Under the generalized linear model with any $\Delta \in (0, 0.5)$, the COUNTING test is guaranteed to control for Type-I error at the level α , and also satisfies the requirement of non-trivial power.*
- (b) *Under the generalized logistic model with any $\tilde{\Delta} > 0$, the DISAGREEMENT test is guaranteed to control for Type-I error at the level α , and also satisfies the requirement of non-trivial power.*

Remark. 1. Result of Theorem 2(a) holds even for a subjective version of the generalized linear model in which for each $(i, j) \in [n] \times [m]$ we substitute q_j with q_{ij} (that is, different values across reviewers) in equations (4a) and (4b), thereby accounting for subjectivity of reviewers.

2. If the logistic model (1) assumed by Tomkins et al. is correct for both SB and DB reviewers with possibly different intercepts, then Theorem 2(b) ensures that the DISAGREEMENT test provably controls for the Type-I error and can detect a bias with probability that goes to 1 as sample size grows, without requiring knowledge (neither exact nor approximate) of papers’ scores q_1^*, \dots, q_n^* .

3. Notice that the COUNTING and DISAGREEMENT tests do not require the knowledge of Δ and $\tilde{\Delta}$ parameters to control for Type-I error and satisfy the requirement of non-trivial power.

Figure 5 compares the performance of the DISAGREEMENT and COUNTING tests under specific instances of the generalized linear and logistic models, illustrating the results of Theorem 2. Figure 5a shows that the COUNTING test controls for Type-I error and has a non-trivial power under the generalized linear model. Notice that the DISAGREEMENT test does not control for Type-I error in this instance, implying that Theorem 2(b) cannot be extended to guarantee the Type-I error control under the generalized linear model as well. Symmetrically, Figure 5b shows that while the DISAGREEMENT test leads to a reliable testing under the generalized logistic model, the COUNTING test is unable to control for Type-I error under this model. Hence, we conclude that under the setup of the relative bias problem, none of the tests dominates another, each leading to a reliable testing under the corresponding model.

We see in Figure 5 that neither the DISAGREEMENT nor the COUNTING test is suitable for *both* generalized linear and logistic models. In the next section we show that this is not a drawback of these specific tests, but rather a manifestation of a more general impossibility result.

6.2.2 Negative result

We conclude our analysis with a negative result that limits the complexity of the class \mathcal{F} for which a reliable testing in the relative bias problem is possible. Let us first state the main result of this section.

Theorem 3. *Suppose that there exist two functions $g, h \in \mathcal{F}$ and some $0 \leq x_1 < x_2 \leq 1$ such that $g(x_1) < h(x_1)$ and $g(x_2) > h(x_2)$. Suppose also that there exists a testing procedure ψ operating on decisions of SB and DB reviewers that for any given $\alpha \in (0, 1)$ keeps Type-I error below α for all matrices Π^{sb}, Π^{db} that satisfy the null hypothesis of Problem 2 specified by some function $f_0 \in \mathcal{F}$. Then the testing procedure ψ cannot satisfy the requirement of non-trivial power.*

The intuition behind Theorem 3 is as follows. If the class \mathcal{F} contains “too many” functions, then some matrices Π^{sb} and Π^{db} satisfy the null hypothesis of Problem 2 defined by $f_0 \in \mathcal{F}$ and simultaneously satisfy

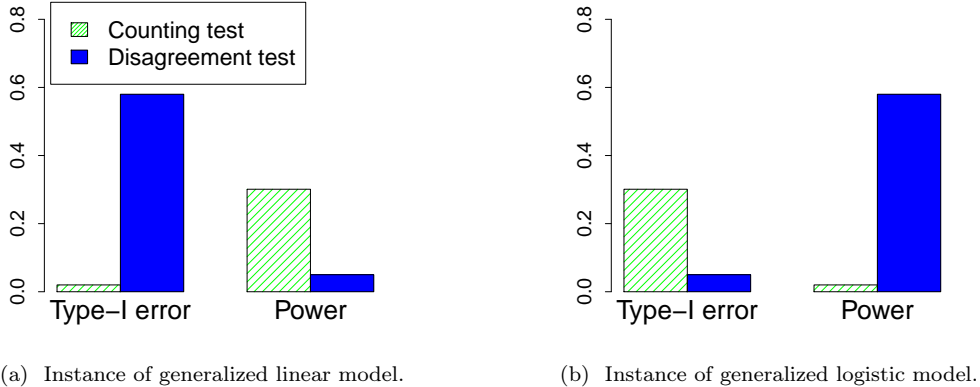


Figure 5: Synthetic simulations evaluating performance of the DISAGREEMENT and COUNTING tests under setup of the relative bias problem when generalized linear (Figure 5a) or generalized logistic (Figure 5b) models is correct. For each model presence and absence of bias are simulated. Details of the simulation setup are provided in Appendix D.3. Error bars are too small to be visible.

the alternative hypothesis (3) defined by some other function $f'_0 \in \mathcal{F}$ with margin $\delta > 0$. Hence, any testing procedure must either have high Type-I error rate or sacrifice the non-trivial power requirement over the class of functions \mathcal{F} , implying that reliable testing is impossible.

Using Theorem 3 we can deduce that one can hope to control for the Type-I error rate and simultaneously have a non-trivial power only when functions contained in \mathcal{F} are pointwise totally ordered, that is, for any two functions $f, g \in \mathcal{F}$ it must be the case that either $f(x) \geq g(x)$ for all $x \in [0, 1]$ or $f(x) \leq g(x)$ for all $x \in [0, 1]$.

Let us now illustrate the consequences of Theorem 3 for the generalized linear and logistic models.

Corollary 1. *For any significance level $\alpha \in (0, 1)$, let ψ_1 and ψ_2 be any testing procedures which operate on decisions of SB and DB reviewers. Suppose that under the generalized linear model with any $\Delta \in (0, 0.5)$, procedure ψ_1 controls for the Type-I error at the level α and satisfies the non-trivial power requirement. Suppose also that under the generalized logistic model with any $\tilde{\Delta} > 0$, procedure ψ_2 controls for the Type-I error at the level α and satisfies the non-trivial power requirement. Then*

- (a) *Under the generalized logistic model with any $\tilde{\Delta} > 0$, procedure ψ_1 incurs a Type-I error rate strictly greater than α .*
- (b) *Under the generalized linear model with any $\Delta \in (0, 0.5)$, procedure ψ_2 incurs a Type-I error rate strictly greater than α .*

Corollary 1 shows that there does not exist a testing procedure that controls over Type-I error and has non-trivial power under both generalized linear and generalized logistic models. As a result, one needs to design different procedures for these models as we did with the DISAGREEMENT and COUNTING tests. In case of the DISAGREEMENT and COUNTING tests, Corollary 1 is illustrated by Figure 5.

In Appendix C we discuss another application of Theorem 3 in context of the generalized logistic model that also suggests that generality of the DISAGREEMENT test cannot be further increased.

7 Discussion

In this work we consider the problem of testing for biases in peer review. We show that under various conditions the approach used by prior work does not control over the Type-I error rate. We underscore that

we do not aim at confirming or disproving the presence of biases found in prior work, but our focus is on the validity of testing methods. With this goal in mind, we propose a principled approach to testing for biases and design two statistical procedures that coupled with our novel experimental setup provably control for the Type-I error rate and have non-trivial power under essentially a single assumption of no difference in behavior of SB and DB reviewers when the bias is absent. We then show that this assumption cannot be relaxed in general and that to accommodate the aforementioned difference in behavior one needs to make some modelling assumptions, as we demonstrated with our tests and generalizations of popular linear and logistic models.

We presented the DISAGREEMENT and COUNTING tests in context of peer review. However, we underscore that one can adapt our experimental setup (Procedure 1) to use our testing procedures in other applications. These applications include peer grading, university admission and hiring where some protected attributes might be available to reviewers.

There are number of open problems suggested by our work. The first direction is associated with statistical power of the testing procedures we propose. In this work we show that our tests have power that going to one under certain conditions on the alternative. It is of interest to establish a bound on the statistical power of our tests in a finite sample setting and compare it with an upper bound on the maximum power that can be achieved by any computationally-efficient testing procedure.

The second direction is related to the design of the experimental procedure. To accommodate tests for biases, one needs to deviate from the standard peer review, thus introducing a trade-off between the quality of the peer-review process and the accuracy of the testing. Quantification of such a trade-off may help to design a better setup and understand the cost of the experiment in terms of the peer-review quality. In this work we designed a procedure that leads to the desired accuracy, but is suboptimal in terms of the TPMS objective. In contrast, the optimal TPMS assignment would not allow to perform the reliable testing. Hence, an open problem is to design an experimental procedure which accommodates our statistical tests and subject to this maximizes the quality of the assignment in terms of the TPMS objective.

References

- Arnold, D., Dobbie, W., and Yang, C. S. (2018). Racial bias in bail decisions*. *The Quarterly Journal of Economics*, 133(4):1885–1932.
- Bakanic, V., McPhail, C., and Simon, R. J. (1987). The manuscript review and decision-making process. *American Sociological Review*, pages 631–642.
- Bertrand, M. and Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American economic review*, 94(4):991–1013.
- Blank, R. M. (1991). The effects of double-blind versus single-blind reviewing: Experimental evidence from the american economic review. *American Economic Review*, 81(5):1041–1067.
- Brunner, J. and Austin, P. C. (2009). Inflation of Type I error rate in multiple regression when independent variables are measured with error. *Canadian Journal of Statistics*, 37(1):33–46.
- Budden, A. E., Tregenza, T., Aarssen, L. W., Koricheva, J., Leimu, R., and Lortie, C. J. (2008). Double-blind review favours increased representation of female authors. *Trends in Ecology and Evolution*, 23(1):4 – 6.
- Charlin, L. and Zemel, R. S. (2013). The Toronto Paper Matching System: An automated paper-reviewer assignment system. In *ICML Workshop on Peer Reviewing and Publishing Models*.
- Ernst, E. and Resch, K.-L. (1994). Reviewer bias: a blinded experimental study. *The Journal of laboratory and clinical medicine*, 124(2):178–182.
- Hill, S. and J. Provost, F. (2003). The myth of the double-blind review? Author identification using only citations. *SIGKDD Explorations*, 5:179–184.

- Kerr, S., Tolliver, J., and Petree, D. (1977). Manuscript characteristics which influence acceptance for management and social science journals. *Academy of Management Journal*, 20(1):132–141.
- Lamont, M. (2009). *How professors think*. Harvard University Press.
- Largent, E. and Snodgrass, R. (2016). Blind peer review by academic journals. In Robertson, C. and Kesselheim, A., editors, *Blinding as a Solution to Bias: Strengthening Biomedical Science, Forensic Science, and Law*, pages 75–95. Cambridge.
- Mahoney, M. J. (1977). Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive therapy and research*, 1(2):161–175.
- Noothigattu, R., Shah, N., and Procaccia, A. (2018). Choosing how to choose papers. *arXiv preprint arxiv:1808.09057*.
- Okike, K., Hug, K. T., Kocher, M. S., and Leopold, S. S. (2016). Single-blind vs double-blind peer review in the setting of author prestige. *JAMA*, 316(12):1315–1316.
- Seeber, M. and Bacchelli, A. (2017). Does single blind peer review hinder newcomers? *Scientometrics*, 113(1):567–585.
- Snodgrass, R. (2006). Single- versus double-blind reviewing: An analysis of the literature. *SIGMOD Record*, 35:8–21.
- Squazzoni, F. and Gandelli, C. (2012). Saint Matthew strikes again: An agent-based model of peer review and the scientific community structure. *Journal of Informetrics*, 6(2):265–275.
- Stefanski, L. A. and Carroll, R. J. (1985). Covariate measurement error in logistic regression. *Ann. Statist.*, 13(4):1335–1351.
- Thorngate, W. and Chowdhury, W. (2014). By the numbers: Track record, flawed reviews, journal space, and the fate of talented authors. In *Advances in Social Simulation*, pages 177–188. Springer.
- Thornhill, T. (2018). We want black students, just not you: How white admissions counselors screen black prospective students. *Sociology of Race and Ethnicity*, page 2332649218792579.
- Tomkins, A., Zhang, M., and Heavlin, W. D. (2017). Reviewer bias in single- versus double-blind peer review. *Proceedings of the National Academy of Sciences*, 114(48):12708–12713.
- Wang, J. and Shah, N. B. (2018). Your 2 is my 1, your 3 is my 9: Handling arbitrary miscalibrations in ratings. *CoRR*, abs/1806.05085.
- Webb, T. J., O’Hara, B., and Freckleton, R. P. (2008). Does double-blind review benefit female authors? *Trends in Ecology and Evolution*, 23(7):351 – 353.
- Weisberg, S. (2005). *Applied Linear Regression*. Wiley, Hoboken NJ, third edition.

Appendix

A More than one property of interest

Throughout the main body of the paper, we considered the case of a single property of interest. We now generalize some of the results to the case of more than one property of interest. First of all, we recall some notation. Let k be the total number of properties, then for each paper $j \in [n]$, we let variables $w_j^{(1)}, w_j^{(2)}, \dots, w_j^{(k)}$ indicate whether or not the paper satisfies the corresponding property. For each property ℓ , the set $\mathcal{J}_\ell \subseteq [n]$ contains papers that satisfy a property ℓ with $\overline{\mathcal{J}}_\ell$ being its complement.

We argue that when $k > 1$, one needs to think about the bias testing problem as of an instance of the relative bias problem as defined in Section 4.2. Indeed, consider for example the case of two properties of interest ($k = 2$) and assume that we are interested in testing for biases with respect to the first property. Then even if there is no bias with respect to this property, the behavior of reviewers between SB and DB conditions might be different due to possible biases with respect to the second property.

The negative result of Theorem 3 we established in Section 6.2.2 also applies to the case of multiple properties, implying that the reliable testing is possible only under some restrictions on the difference in reviewers' behavior between SB and DB conditions under the absence of bias. Following the relative bias problem defined as Problem 2, we now generalize it for $k > 1$. To this end, we consider a problem of testing for biases with respect to the property $\ell \in [k]$ and introduce an additional piece of notation. For each paper $j \in [n]$, let \mathbf{w}_j denote a vector of indicators of property satisfaction: $\mathbf{w}_j = (w_j^{(1)}, w_j^{(2)}, \dots, w_j^{(k)})$ and let $\mathbf{w}_j^{(-\ell)}$ denote the same vector but with ℓ^{th} component omitted, that is, $\mathbf{w}_j^{(-\ell)} = (w_j^{(1)}, \dots, w_j^{(\ell-1)}, w_j^{(\ell+1)}, \dots, w_j^{(k)})$.

Following the definition of the relative bias problem (Problem 2), the set \mathcal{F}_ℓ contains functions that under the absence of bias with respect to the property ℓ , specify the difference in behavior between DB and SB conditions. In case of a single property of interest, \mathcal{F} was a subset of all non-decreasing functions $f : [0, 1] \rightarrow [0, 1]$. However, when $k > 1$, even under the absence of the bias with respect to the property ℓ , the change in reviewers' behavior between DB and SB conditions may be influenced by whether the paper satisfies properties other than ℓ , due to possible biases with respect to these properties. Hence, under the absence of bias with respect to the property ℓ , the change of behavior between SB and DB conditions is described as follows:

$$\forall (i, j) \in [n] \times [m] : \pi_{ij}^{(\text{sb})} = f_0(\pi_{ij}^{(\text{db})}, \mathbf{w}_j^{(-\ell)}),$$

where function $f_0 \in \mathcal{F}_\ell$ is non-decreasing in its first argument. Thus, the set \mathcal{F}_ℓ is a subset of all functions with domain $[0, 1] \times \{0, 1\}^{k-1}$ which are non-decreasing in their first argument, that is,

$$\mathcal{F}_\ell \subseteq \left\{ f : [0, 1] \times \{0, 1\}^{k-1} \rightarrow [0, 1] \mid f \text{ is non-decreasing in its first argument} \right\}.$$

Having defined the necessary notation, we are ready to introduce the relative bias problem in case of multiple properties of interest.

Problem 3 (Relative bias problem for multiple properties.). Given significance level $\alpha \in (0, 1)$, the property of interest ℓ , the class of functions \mathcal{F}_ℓ and decisions of SB and DB reviewers, the goal is to test the following hypotheses:

$$\begin{aligned} H_0 : \forall i \in [m] \forall j \in [n] \quad \pi_{ij}^{(\text{sb})} &= f_0(\pi_{ij}^{(\text{db})}, \mathbf{w}_j^{(-\ell)}) \\ H_1 : \forall i \in [m] \forall j \in [n] \quad \begin{cases} \pi_{ij}^{(\text{sb})} \geq f_0(\pi_{ij}^{(\text{db})}, \mathbf{w}_j^{(-\ell)}) & \text{if } j \in \mathcal{J}_\ell \\ \pi_{ij}^{(\text{sb})} \leq f_0(\pi_{ij}^{(\text{db})}, \mathbf{w}_j^{(-\ell)}) & \text{if } j \notin \mathcal{J}_\ell \end{cases} \end{aligned} \quad (10)$$

for some unknown $f_0 \in \mathcal{F}_\ell$, and where at least one inequality in the alternative hypothesis (10) is strict.

Intuitively, under the null hypothesis of absence of bias with respect to the property ℓ , the change of reviewers' behaviour between SB and DB conditions is determined by (i) bias introduced by SB condition itself which is independent of papers' authorship information and (ii) bias with respect to properties other than the property ℓ . The generalized linear and logistic models can be formulated in case of multiple properties of interest as follows.

Generalized linear model. Given a fixed constant $\Delta \in (0, 0.5)$, we follow the case of a single property and assume that each paper $j \in [n]$ has some unknown representation $q_j \in (\Delta, 1 - \Delta)$. The generalized linear model assumes that for each $(i, j) \in [m] \times [n]$, the behaviour of reviewer i if she/he reviews paper j is described by the following parametric equations:

$$\text{DB: } \pi_{ij}^{(\text{db})} = q_j \tag{11a}$$

$$\text{SB: } \pi_{ij}^{(\text{sb})} = q_j + \beta_0^{(\text{sb})} + \sum_{\ell \in [k]} \beta_\ell^{(\text{sb})} w_j^{(\ell)}, \tag{11b}$$

where unknown coefficients are such that $|\beta_0^{(\text{sb})}| + \sum_{\ell=1}^k |\beta_\ell^{(\text{sb})}| < \Delta$. Under the generalized linear model, a bias with respect to the property ℓ is present whenever $\beta_\ell^{(\text{sb})} \neq 0$.

Generalized logistic model. Given a fixed constant $\tilde{\Delta} > 0$, the generalized logistic model assumes that (i) for every $j \in [n]$, a corresponding representation q_j belongs to the interval $(-\tilde{\Delta}, \tilde{\Delta})$ and (ii) for each $(i, j) \in [m] \times [n]$, the behaviour of reviewer i if she/he reviews paper j is described by the following parametric equations:

$$\text{DB: } \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} = \beta_0^{(\text{db})} + \beta_1 q_j \tag{12a}$$

$$\text{SB: } \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = \beta_0^{(\text{sb})} + \beta_1 q_j + \sum_{\ell \in [k]} \beta_{\ell+1}^{(\text{sb})} w_j^{(\ell)}, \tag{12b}$$

where all coefficients are bounded in absolute value by $\tilde{\Delta}$ and $\beta_1 > 0$. Under the generalized logistic model, a bias with respect to the property ℓ is present whenever $\beta_{\ell+1}^{(\text{sb})} \neq 0$.

Remark. 1. First, provided that matrix Π^{db} is generated according to the one of the introduced models, one can define a set of functions \mathcal{F}_ℓ that puts the corresponding model in the context of the relative bias problem for multiple properties of interest defined as Problem 3.

2. The goal under each of the models introduced above is to test the significance of the coefficient in equation describing the behavior of SB reviewer that corresponds to the indicator of the property of interest. For example, if we are interested in testing for biases with respect to the property ℓ under the generalized logistic model, then we want to test the significance of the coefficient $\beta_{\ell+1}^{(\text{sb})}$ in equation (12b).

3. Notice that in case of multiple properties, the models we introduced above describe reviewers' behaviour both under the absence of bias and under the presence of bias. In this way they allow simultaneous testing for biases with respect to many properties of interest.

Observe that the relationships that describe the behaviour of SB reviewers in models (11b) and (12b) are reminiscent of the linear regression and logistic regression models respectively. As mentioned above, one cannot fit decisions of SB reviewers to these models using existing methods, because one of the covariates (paper representation q) is unknown. In their work, Tomkins et al. employed DB reviewers to estimate this unknown covariate and used these estimates to fit the logistic model. As we discussed in Section 3, this approach leads to an unreliable testing procedure under various realistic conditions.

We now show that using ideas of the DISAGREEMENT and COUNTING tests, one can use decisions of both SB and DB reviewers to eliminate the unknown covariate from the model, thereby enabling standard tools without the need to estimate any covariate.

Proposition 1. *Let reviewers i and i' be assigned to paper j in SB and DB setups correspondingly, then*

- (a) *Under the generalized linear model the expectation of the quantity $Y_{ij} - X_{i'j}$ follows the linear model (11b) with $q_j = 0$:*

$$\mathbb{E}[Y_{ij} - X_{i'j}] = \beta_0^{(sb)} + \sum_{\ell \in [k]} \beta_\ell^{(sb)} w_j^{(\ell)}. \quad (13a)$$

- (b) *Under the generalized logistic model the expectation of the quantity $Y_{ij} | (Y_{ij} \neq X_{i'j})$ follows the logistic model (12b) with $q_j = 0$:*

$$\log \frac{\mathbb{E}[Y_{ij} | (Y_{ij} \neq X_{i'j})]}{1 - \mathbb{E}[Y_{ij} | (Y_{ij} \neq X_{i'j})]} = \beta_0 + \sum_{\ell \in [k]} \beta_{\ell+1}^{(sb)} w_j^{(\ell)}, \quad (13b)$$

where $\beta_0 = \beta_0^{(sb)} - \beta_0^{(db)}$.

Proposition 1 provides a mean to eliminate the unknown covariate from the models of SB decisions (11b) and (12b) using decisions of DB reviewers. For example, in case of the generalized logistic model it relies on the core idea of the DISAGREEMENT test and suggests conditioning on pairs (SB reviewer, DB reviewer) such that reviewers disagree in their decisions for some paper. After conditioning, decisions of SB reviewers follow model (13b) with all covariates known and hence standard test for logistic regression can be applied to evaluate significance of coefficients.

Proposition 1 also allows to avoid using noisy measurements and hence any test for significance of the coefficients applied to the models (13a) and (13b) will not be susceptible to issues caused by the use of noisy measurements and misspecification of the meaning of q (issues (a) and (b) from Section 3). If one restricts each reviewer to input at most one decision to the testing procedure, then issue (c) will also be mitigated.

B Our tests under the setup of Tomkins et al.

In this section we give additional comments on the applicability of our testing procedures to the setup of Tomkins et al. (2017). To this end, recall that our tests take as input the set of tuples \mathcal{T} such that (i) each tuple $t \in \mathcal{T}$ is of the form $t = (j_t, Y_{j_t}, X_{j_t}, w_{j_t})$, where j_t is a corresponding paper, Y_{j_t}, X_{j_t} are decisions of SB and DB reviewers for this paper and w_{j_t} equals 1 if $j_t \in \mathcal{J}$ and -1 otherwise; (ii) each reviewer contributes at most one decision to the set \mathcal{T} . Potentially our tests can be coupled with any experimental procedure as long as this procedure enables a construction of such a set \mathcal{T} . However, one needs to understand that while Procedure 1 is robust to issues we discussed in Section 3.2, other experimental setups may lead to an inflation of the Type-I error.

In the experiment conducted by Tomkins et al. (2017), reviewers were split into two groups (SB and DB) uniformly at random at the very beginning of the experiment. Then two assignments A_{SB} and A_{DB} were computed separately for each group of reviewers. As discussed in Section 3.2, even if both groups of reviewers bid in a blind manner, the design of Tomkins et al. may lead to inflated Type-I error. Nonetheless, we now show that even under the setup of Tomkins et al., one can employ the DISAGREEMENT and COUNTING tests to fix issues (a)-(c) with the testing procedure discussed in Section 3.1.

B.1 Matching algorithms

Let us first introduce two matching procedures that construct an input for our tests under the setup of Tomkins et al. (2017). Given assignment of papers to reviewers in both SB and DB conditions, we discuss two choices of matching algorithms depending on the relationship between parameters λ (required number of reviewers per paper in each condition) and μ (maximum number of papers per reviewer). Notice that our goal is not to maximize the size of \mathcal{T} , but instead to maximize the minimum of the number of papers from \mathcal{J} included in

the set \mathcal{T} and the number of papers from $\overline{\mathcal{J}}$ included in the set \mathcal{T} . This is because our statistical tests need decisions for papers from both \mathcal{J} and $\overline{\mathcal{J}}$ to maximize their power. Depending on the relationship between λ and μ we can solve this problem either exactly or approximately.

Case 1 ($\lambda \geq \mu$). In this case, each paper can be matched to 1 SB reviewer and 1 DB reviewer by finding two separate maximum matchings (papers to SB reviewers and papers to DB reviewers) using the Hungarian matching algorithm. We formally present the matching procedure as Algorithm 1.

Algorithm 1 Exact matching algorithm

Input: Assignments A_{SB} , A_{DB} of SB and DB reviewers to papers, respectively.

1. Construct a graph G that consists of 3 layers:

- **Layer 1.** One node for each SB reviewer
- **Layer 2.** One node for each paper
- **Layer 3.** One node for each DB reviewer

and add edges between reviewers and papers according to assignments A_{SB} and A_{DB} . Set $\mathcal{T} = \emptyset$.

2. Using the Hungarian matching algorithm with uniform tie-breaking find matchings \mathcal{M}_{SB} and \mathcal{M}_{DB} where \mathcal{M}_{SB} (respectively \mathcal{M}_{DB}) is a maximum 1-1 matching between SB (respectively DB) reviewers and papers (each reviewer is matched to at most 1 paper and each paper is matched to at most 1 reviewer).

3. Leave in graph G only those edges that correspond to matched pairs in \mathcal{M}_{SB} and \mathcal{M}_{DB} .

4. For any triple of (SB reviewer i_1 , paper j , DB reviewer i_2) such that there is a path from a node that corresponds to reviewer i_1 to a node that corresponds to reviewer i_2 through a node that corresponds to paper j , add $t = (j, Y_{i_1 j}, X_{i_2 j}, w_j)$ to \mathcal{T} .

5. Return \mathcal{T} .

Lemma 1. *For any assignments of SB and DB referees to papers that satisfy (λ, μ) -load constraints with $\lambda \geq \mu$, the matching procedure in Algorithm 1 is guaranteed to construct a set of tuples \mathcal{T} such that for each paper $j \in [n]$ there is one tuple that corresponds to this paper.*

Case 2 ($\lambda < \mu$). In this case we cannot use the above idea, because there does not exist a matching such that each paper is matched to one SB and one DB reviewer, subject to a constraint that each reviewer is matched with at most one paper. While solving the exact optimization problem in this case might be hard, a simple greedy procedure constructs a sufficiently large matching for the DISAGREEMENT and COUNTING tests to satisfy the non-trivial power requirement. The iterative greedy procedure in each iteration matches one paper from $\overline{\mathcal{J}}$ and one paper from \mathcal{J} to 1 SB and 1 DB reviewer and removes those reviewers from subsequent iterations to maintain the constraint that each reviewer contributes at most one decision to the set \mathcal{T} . We formally introduce the greedy procedure as Algorithm 2.

Lemma 2. *For any assignments of SB and DB referees to papers that satisfy (λ, μ) -load constraints, the matching procedure in Algorithm 2 is guaranteed to construct a set of tuples \mathcal{T} that for large enough $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ contains at least $c \min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ tuples corresponding to papers from \mathcal{J} and at least $c \min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ tuples corresponding to papers from $\overline{\mathcal{J}}$, where c is a constant that may depend only on λ and μ .*

Remark. 1. If the set \mathcal{T} constructed by the Algorithm 1 is such that there exist reviewers who do not contribute any of their decisions to this set, then one can run Algorithm 2 on assignments of these reviewers to papers and obtain the set \mathcal{T}' . Next, consider the updated set $\mathcal{T}^* = \mathcal{T} \cup \mathcal{T}'$ and observe that each reviewer contributes at most one decision to this set.

2. By construction both matching algorithms introduced in this section include at most one decision per reviewer in a set of tuples \mathcal{T} .

Overall, let \mathcal{A} denote a procedure that takes assignments A_{SB} and A_{DB} as input and depending on the relationship between λ and μ calls Algorithm 1 or Algorithm 2 to construct the set \mathcal{T} .

Algorithm 2 Greedy matching algorithm

Input: Assignments A_{SB} , A_{DB} of SB and DB reviewers to papers, respectively.

1. Construct a graph G that consists of 3 layers:

- **Layer 1.** One node for each SB reviewer
- **Layer 2.** One node for each paper
- **Layer 3.** One node for each DB reviewer

and add edges between reviewers and papers according to assignments A_{SB} and A_{DB} . Set $\mathcal{T} = \emptyset$.

2. Find a triple (SB reviewer i_1 , paper $j \in \mathcal{J}$, DB reviewer i_2) such that there is path in graph G from a node corresponding to SB reviewer to a node corresponding to DB reviewer through a node corresponding to a paper. If there are many such triples, break ties uniformly at random. If such a triple exists, define $t_1 = (j, Y_{i_1 j}, X_{i_2 j}, w_j)$, otherwise set $t_1 = \emptyset$.

3. Find a triple (SB reviewer $i'_1 \neq i_1$, paper $j' \in \overline{\mathcal{J}}$, DB reviewer $i'_2 \neq i_2$) such that there is path in graph G from a node corresponding to the SB reviewer to a node corresponding to the DB reviewer through a node corresponding to the paper. If there are many such triples, break ties uniformly at random. If such a triple exists, define $t_2 = (j', Y_{i'_1 j'}, X_{i'_2 j'}, w_{j'})$, otherwise set $t_2 = \emptyset$.

4. Update $\mathcal{T} = \mathcal{T} \cup \{t_1, t_2\}$. If both t_1 and t_2 are empty, return \mathcal{T} . Otherwise delete reviewers i_1, i'_1, i_2, i'_2 from the graph G together with the corresponding edges and go to Step 2.

B.2 Guarantees

Having defined a procedure to construct input for the DISAGREEMENT and COUNTING tests, we are now ready to formulate corresponding theoretical guarantees. Recalling that the experimental setup of Tomkins et al. itself breaks the Type-I error guarantees, we abstract out these issues by assuming that instead of TPMS assignment algorithm or any other algorithm that computes assignments of papers to SB and DB reviewers, the assignment is selected uniformly at random from the set of all assignments satisfying (λ, μ) -constraints. For brevity, we only show the result for the absolute bias problem, but analogue of Theorem 2 also holds.

Proposition 2. *For any given significance level $\alpha \in (0, 1)$, under the setup of the absolute bias problem (Problem 1), let experiment be organized according to the procedure of Tomkins et al. with random assignment. Then the DISAGREEMENT and COUNTING tests coupled with procedure \mathcal{A} (Algorithm 1 and Algorithm 2) are guaranteed to control for the Type-I error rate at the level α and also satisfy the requirement of non-trivial power.*

Remark. 1. From theoretical standpoint, the requirement of random assignment can be substituted with the following conditions, under which any assignment algorithm can be used: (i) reviewers in both conditions bid blindly and (ii) reviewers' evaluations are independent of similarities. That is, rows of matrices Π^{sb} and Π^{db} are assigned to reviewers uniformly at random after the assignment is computed.

2. Proposition 2 ensures that the COUNTING and DISAGREEMENT tests, coupled with matching algorithms we introduced above, are robust to issues (a)-(c) discussed in Section 3.1. However, in practice even our robust tests may still be susceptible to issues caused by the experimental setup of Tomkins et al.

C Additional impossibility result for the generalized logistic model

In this section we formulate an additional impossibility result that highlights the generality of the DISAGREEMENT test.

As we demonstrated in Theorem 2, the DISAGREEMENT test leads to a reliable testing under the generalized logistic model. Recalling the definition of papers representations $q_j, j \in [n]$, let us now consider an extended version of the generalized logistic model which is given by the following parametric equations describing the behaviour of DB and SB reviewers under the absence of bias

$$\text{DB: } \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} = \beta_0^{(\text{db})} + \beta_1^{(\text{db})} q_j \quad (14a)$$

$$\text{SB: } \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = \beta_0^{(\text{sb})} + \beta_1^{(\text{sb})} q_j, \quad (14b)$$

where parameters $\beta_0^{(\text{sb})}, \beta_1^{(\text{sb})} > 0, \beta_0^{(\text{db})}, \beta_1^{(\text{db})} > 0$ as well as papers' scores $q_j, j \in [n]$, are bounded in absolute value but otherwise are allowed to be arbitrary. This extended version specializes to the standard generalized logistic model if $\beta_1^{(\text{sb})} = \beta_1^{(\text{db})}$. In words, under the extended version, the reviewers in SB and DB conditions under the absence of bias may have not only different intercepts, but also different coefficients in front of q . The presence of bias is then defined as a violation of model (14b) where the direction of violation is different for papers from \mathcal{J} and $\bar{\mathcal{J}}$.

Unfortunately, as we show in Corollary 2, the reliable testing under this extension of the generalized logistic model is not possible.

Corollary 2. *For any significance level $\alpha \in (0, 1)$, consider the extension of the generalized logistic model given by equations (14a) and (14b). Then no test operating on decisions of SB and DB reviewers can control for the Type-I error rate at the level α and simultaneously satisfy the non-trivial power requirement.*

The negative result of Corollary 2 applies to the DISAGREEMENT test, implying that even our testing procedure which is robust to various issues discussed in Section 3 cannot handle the extension of the generalized logistic model specified by equations (14a) and (14b).

D Setup for simulations

In this section we describe setup for simulations we conducted in this work. Notice that in contrast to the test of Tomkins et al. (2017) which operates on accept/reject decisions of SB reviewers and scores provided by DB reviewers, the tests we introduce in this work operate on decisions of both SB and DB reviewers. Hence, to compare tests we need to specify (i) models of DB/SB reviewers' decisions and (ii) models of DB reviewers' scores. All simulations are run for 5000 iterations.

D.1 Simulations in Section 3

We now provide necessary details for the simulations in Section 3.

D.1.1 Measurement error (Figure 2a)

For this simulation we consider the following model of SB and DB reviewers:

$$\text{DB: } \log \frac{\pi_j^{(\text{db})}}{1 - \pi_j^{(\text{db})}} = \beta_0 + \beta_1 q_j^* \quad (15a)$$

$$\text{SB: } \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = \beta_0 + \beta_1 q_j^* + \beta_2 w_j, \quad (15b)$$

that is, model (1) is correct and reviews given by the same reviewer for different papers are independent. Notice that under this model all reviewers are identical and hence issues with the setup do not manifest in this case.

We set $m = 2n = 1000$ and $\mu = \lambda = 2$. At each iteration we independently sample true scores of papers $q_j^*, j \in [n]$, from uniform distribution $\mathcal{U}[-2, 2]$ and assume that mean scores by two DB reviewer assigned to a paper $j \in [n]$ estimates true score q_j^* with some Gaussian noise ($\sigma = 0.7$). We then sample values of $w_j, j \in [n]$, such that correlation between q^* and w equals φ for values of φ between 0 and 0.5. To this end, we let each paper $j \in [n]$ with the score $q_j^* < 0$ have $w_j = 1$ with probability $0.5 - \gamma$ and $w_j = -1$ otherwise. Similarly, each paper $j \in [n]$ with the score $q_j^* \geq 0$ has $w_j = 1$ with probability $0.5 + \gamma$ and $w_j = -1$ otherwise. We then vary the value of $\gamma \in (0, 0.5)$ to achieve the necessary correlation. Finally, using models (15a) and (15b) with $\beta_0 = 1, \beta_1 = 2$ and $\beta_2 = 0$ (no bias condition) we sample decisions of SB and DB reviewers and run the DISAGREEMENT test and the test used by Tomkins et al. (2017), setting the significance level to be $\alpha = 0.05$. We then compute a Type-I error as a fraction of iterations in which the null hypothesis ($\beta_2 = 0$) was rejected.

D.1.2 Model mismatch (Figure 2b)

For this simulation we consider a violation of model (1) and the following model of SB and DB reviewers with $\beta_2 = 0$ (no-bias condition):

$$\begin{aligned} \text{DB: } \log \frac{\pi_j^{(\text{db})}}{1 - \pi_j^{(\text{db})}} &= \beta_0 + \beta_1 (q_j^*)^3 \\ \text{SB: } \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} &= \beta_0 + \beta_1 (q_j^*)^3 + \beta_2 w_j. \end{aligned}$$

To abstract out the effect of measurement error, in this section we assume that the true scores $q_j^*, j \in [n]$, are known, but the test used by Tomkins et al. (2017) fits the model defined by equation (15b). Besides the change of correct model and availability of true scores $\{q_j^*, j \in [n]\}$, the simulations follow scenario we described in Appendix D.1.1.

D.1.3 Reviewer calibration (Figure 2c)

In this simulation we model the effect of correlations introduced by reviewer calibration. More concretely, we construct a model of reviewer calibration under which the test by Tomkins et al. (2017) fails to control for the Type-I error rate. In this section we assume that true scores of submissions are proportional to the clarity of the writing. We then sample clarity scores $\zeta_j, j \in [n]$, from uniform distribution $\mathcal{U}[-1, 1]$ and define $q_j^* = \zeta_j$ for each $j \in [n]$. Eventually, we consider the following model of reviewer. For each $i \in [m]$ and for each $j \in [n]$:

$$\begin{aligned} \text{DB: } \pi_{ij}^{(\text{db})} &= \pi_j^{(\text{db})} + \ell_i \times \mathbb{I}[\zeta_j < 0.5] \\ \text{SB: } \pi_{ij}^{(\text{sb})} &= \pi_j^{(\text{db})} + \ell_i \times \mathbb{I}[\zeta_j < 0.5], \end{aligned}$$

where ℓ_i is reviewers' leniency which equals 0.4 with probability 0.5 and -0.4 otherwise and $\pi_j^{(\text{db})}, \pi_j^{(\text{sb})}$ are defined by equations (15a) and (15b) with $\beta_0 = 0, \beta_1 = 0.25$ and $\beta_2 = 0$ (no bias condition). Parameters are selected to ensure that $0 \leq \pi_{ij}^{(\text{db})}, \pi_{ij}^{(\text{sb})} \leq 1$.

In words, the above model says that for papers with high quality of writing ($\zeta > 0.5$) reviewers understand their content well and follow models (15a) and (15b) exactly, but for papers with lower writing quality their leniency parameter influences their decision. Notice that under this model it is natural to expect that estimates of the true scores provided by DB reviewers are also influenced by their leniency and hence are noisy. However, to isolate the effect of reviewer identity we assume that the test used by Tomkins et al. (2017) knows true scores $q_j^*, j \in [n]$, exactly. Additionally, notice that marginally each reviewer follows the model defined by equations (15a) and (15b), and hence when $\mu = 1$, the test by Tomkins et al. (2017) has control over the Type-I error for any correlation between q^* and w .

In this section we consider an extreme pattern of correlations between q^* and w . Concretely, we assume that for any paper $j \in [n]$, we have $w_j = 1$ if and only if $q_j^* > 0.5$ and $w_j = -1$ otherwise. Notice that in practice such strong dependence is unlikely to happen, but we underscore that in practice the test by Tomkins et al. (2017) also does not have access to noiseless true scores which will cause measurement errors and hence will exacerbate the issue.

We then perform simulations as discussed above having $n = 1000$ and $\lambda = 1$ fixed and varying the number of papers per reviewer and using the modification of the Wald test with factor variable for each reviewer added (reviewer-dependent intercept).

D.1.4 Non-blind bidding (Figure 3a)

Formalizing the intuition we mentioned in Section 3.2, we consider a setting with $n = 1000$, $m = 2000$, $\lambda = \mu = 1$ and consider a property of interest “paper has a famous author”. Suppose that during the bidding procedure each reviewer $i \in [m]$ gives a score $b_{ij} \in \{-1, 0, 1\}$ to each paper $j \in [n]$, where $b_{ij} = 1$ means that reviewer wants to review the paper, $b_{ij} = -1$ means that reviewer does not want to review the paper and $b_{ij} = 0$ is an intermediate between $b_{ij} = 1$ and $b_{ij} = -1$. Given the bids, the assignment is computed maximizing the total sum of the bids. Namely, for all $(i, j) \in [m] \times [n]$ let a binary indicator A_{ij} equal 1 if reviewer i is assigned to paper j and 0 otherwise and let $\mathcal{R}_{\text{SB}} \subset [m]$ be the set of reviewers allocated to SB condition. Then the assignment of SB reviewers to papers is computed maximizing the following objective subject to the standard (λ, μ) -load constraints.

$$\sum_{i \in \mathcal{R}_{\text{SB}}} \sum_{j \in [n]} A_{ij} b_{ij}.$$

The same objective is used to assign DB reviewers to papers. Next, we suppose that for each paper $j \in [n]$ there is a true score $q_j^* \in [0, 0.9]$ and that all reviewers belong to one of the following personality types:

- **Type A:** Lenient reviewers who accept each paper $j \in [n]$ assigned to them with probability $q_j^* + 0.1$ and want to read papers from top authors. If bidding is blind, they do not have any information about author identity and bid 0 on each paper, but if bidding is non-blind, then for each paper $j \in \mathcal{J}$ reviewer i of type A places a bid $b_{ij} = 1$ and for each paper $j \in \overline{\mathcal{J}}$ she/he places a bid $b_{ij} = -1$.
- **Type B:** Accurate reviewers who accept each paper $j \in [n]$ assigned to them with probability q_j^* and do not mind reviewing any paper. Independent of whether bidding is blind or not, reviewer i of type B places a bid $b_{ij} = 0$ on each paper $j \in [n]$.

Notice that evaluations of reviewers of both types are unbiased — the probability of acceptance is not determined by author identities. The type of each reviewer is determined independently: reviewer $i \in [m]$ is of type A with probability 0.3 and of type B with probability 0.7. Independently, each paper $j \in [n]$ belongs to \mathcal{J} with probability 0.3 and to $\overline{\mathcal{J}}$ with probability 0.7.

Having defined the setup, in each iteration we independently sample true scores of submissions from $\mathcal{U}[0, 0.9]$ (no correlation with indicator w) and compute two bidding matrices: (i) when SB reviewers observe author identities during bidding and (ii) when bidding is blind for both SB and DB reviewers. For each bidding matrix we compute assignments of SB and DB reviewers to papers and pass observed decisions to the DISAGREEMENT test and the test used by Tomkins et al. (2017). For the test of Tomkins et al., we assume that true scores $q_j^*, j \in [n]$, are known exactly.

D.1.5 Non-random assignment (Figure 3b)

In this section we construct a similarity matrix S and formalize the dependence of reviewer’s perception of a paper on similarity between paper and reviewer that leads to the effect demonstrated in Figure 3b. We notice that the construction we provide here is artificial and serves as a proof of concept for our claim that non-random assignment may violate some key independence assumptions of statistical tests even if it is not

based on reviewers' bids. While in practice we do not expect to observe such specific similarity matrices, we can still observe some more subtle manifestations of issues caused by non-randomness of the assignment.

First, in this section we assume that assignment is performed using the TPMS algorithms (Charlin and Zemel, 2013), that is, given similarity matrix S between reviewers and papers, each paper is assigned to λ reviewers in a way that each reviewer is assigned to at most μ papers such that total sum similarity of the assignment is maximized.

Second, consider a similarity matrix S , defined as follows. For each reviewer $i \in [m]$ and for each paper $j \in [n]$:

$$S_{ij} = (m + 1 - i) \times (n + 1 - j). \quad (16)$$

Given that reviewers are allocated to conditions at random, similarity matrices S_{SB} (SB condition) and S_{DB} (DB condition) are constructed by random division of rows of S into two groups of equal size and stacking them into S_{SB} and S_{DB} correspondingly.

Third, we assume that each reviewer $i \in [m]$ has some value of threshold z_i such that if reviewer i is assigned to paper $j \in [n]$ in either of setups, reviewer accepts the paper with probability π_{ij} given by:

$$\pi_{ij} = \begin{cases} 0.9 & \text{if } S_{ij} \geq z_i \\ q_j^* & \text{if } S_{ij} < z_i, \end{cases} \quad (17)$$

where $q_j^* \in [0, 0.9]$ is a true score of paper j . We also assume that reviewer i in DB condition returns π_{ij} as an estimate of q_j^* .

Fourth, for every reviewer i we set a value of threshold as follows:

$$z_i = (m + 1 - i) \times (n - [(i-1)/2]), \quad (18)$$

where $[x]$ is the integral part of x .

Fifth and finally, we assume that true scores q^* are independently sampled from $\mathcal{U}[0, 0.9]$ and sample indicators w such that they are correlated with q^* , fixing the value of correlation $\varphi = 0.45$. We also set $\mu = \lambda = 1$ and $m = 2n = 1000$. Now we allocate half of reviewers to SB condition and half to DB condition uniformly at random. We then compare the performance of the DISAGREEMENT test and the test by Tomkins et al. under (i) experimental setup of Tomkins et al. and (ii) our experimental setup.

The intuition behind our construction of matrix S in equation (16) is that for any square submatrix of S , the TPMS algorithm with parameters $\mu = \lambda = 1$ will compute an assignment that corresponds to the diagonal of this submatrix. Coupled with specific choice of thresholds (18), probabilities of acceptance (17) and correlation between q^* and w at the level of 0.45, this choice of similarity matrix ensures that under the setup of Tomkins et al., with non-zero probability most of SB reviewers will receive papers with similarities above the corresponding threshold and most of DB reviewers will receive papers with similarities below the corresponding threshold or vice versa. Hence, the assignments will be structurally different and, as demonstrated by Figure 3b, this difference will be confused with bias by both Tomkins et al. (2017) and DISAGREEMENT tests. In contrast, under our proposed setups the assignments of SB and DB reviewers to paper do not exhibit any structural difference and hence do not break the Type-I error guarantees of the tests.

D.2 Simulations in Section 1 and Section 6.1

The simulations in Section 1 and Section 6.1 were performed under the model of reviewers in (15a) and (15b) following the setup described in Appendix D.1.1 with small differences. Instead of varying the value of correlation φ between q^* and w , we fix the value of φ and vary the number of papers n . Moreover, we independently assign papers to the sets \mathcal{J} and $\overline{\mathcal{J}}$ as follows: each paper j such that $q_j^* < 0$ belongs to the set \mathcal{J} with probability $0.5 - \gamma$ and otherwise belongs to the set $\overline{\mathcal{J}}$, similarly, each paper j with $q_j^* > 0$ belongs to the set \mathcal{J} with probability $0.5 + \gamma$ and otherwise belongs to the set $\overline{\mathcal{J}}$. The value of γ is selected to achieve the required level of correlation φ between q^* and w .

- For Figure 1a we set $\varphi = 0.4$ and perform simulations under $\beta_0 = 1, \beta_1 = 2, \beta_2 = 0$ (no bias), $\lambda = 2, \mu = 1$, where true scores are sampled from $\mathcal{U}[-1, 1]$. We see that for the test used by Tomkins et al. (2017) a violation of Type-I error guarantees caused by measurement error coupled with correlations (see Appendix D.1.1 for details) exacerbates as sample size grows.
- For Figure 1b we set $\varphi = 0.6$ and perform simulations under $\beta_0 = 1, \beta_1 = 2, \beta_2 = -0.35$ (bias against papers that satisfy the property), $\lambda = 2, \mu = 1$, where true scores are sampled from $\mathcal{U}[-0.5, 0.5]$. We see that in this case measurement error has strong harmful impact on the power of the test used by Tomkins et al. (2017).
- For Figure 1c we set $\varphi = 0$ and additionally assume that DB reviewers estimate true scores with no noise. In this case all parametric assumptions made by Tomkins et al. (2017) are satisfied. We then perform simulations under $\beta_0 = 1, \beta_1 = 2, \beta_2 = 0.35$ (bias in favour of papers that satisfy the property), $\lambda = 2, \mu = 1$, where true scores are sampled from $\mathcal{U}[-1, 1]$.
- Simulations in Section 6.1 follow the simulations in Figure 1b and Figure 1c with the exception that the COUNTING test is added for comparison.

D.3 Simulations in Section 6.2

In this section we illustrate that the COUNTING test designed to control for Type-I error under the generalized linear model does not lead to a reliable testing under the generalized logistic model under which the DISAGREEMENT test is suitable, and vice versa. To this end, we design two instances of the relative bias problem under the generalized *linear* model — instance (i) with presence of bias and instance (ii) with absence of bias. Our construction ensures that the resulting matrices Π^{sb} and Π^{db} *simultaneously also* fall in the relative bias problem under the generalized *logistic* model with the exception that instance (i) corresponds to absence of bias under the generalized logistic model and instance (ii) corresponds to the presence of bias under this model.

Instance (i) Under the generalized linear model, for each paper $j \in \mathcal{J}$ let $q_j^* = 0.7$ and for each paper $j \in \overline{\mathcal{J}}$ let $q_j^* = 0.5$. Additionally, let $\nu = 0.175$. This choice of parameters defines matrices Π_1^{sb} and Π_1^{db} that are generated according to the equations (4a) and (4b) and fall under the null hypothesis of no bias.

Instance (ii) Under the generalized linear model, for each paper $j \in \mathcal{J}$ let $q_j^* = 0.65$ and for each paper $j \in \overline{\mathcal{J}}$ let $q_j^* = 0.25$. Now let matrix Π_2^{db} be defined according to the model (4a) and matrix Π_2^{sb} be defined as follows:

$$\pi_{ij}^{(\text{sb})} = \begin{cases} q_j^* + \nu_1 & \text{if } w_j = 1 \\ q_j^* + \nu_2 & \text{if } w_j = -1 \end{cases}, \quad (19)$$

where we carefully select $\nu_2 > \nu_1$ as explained below. This choice of parameters leads to a correct alternative hypothesis of presence of bias against papers that satisfy the property of interest.

We now simulate reviewers decisions with $\lambda = 2, \mu = 1, n = 1000, m = 4000$, independently allocating each paper to \mathcal{J} with probability 0.5 and to $\overline{\mathcal{J}}$ otherwise. We then apply the COUNTING and DISAGREEMENT tests for each of these instances, and present the results in Figure 5a. Instance (i) allows to compare Type-I error rates, and instance (ii) allows to compare powers of the tests under the generalized linear model.

One can verify that the instances we constructed above under the generalized *linear* model also fall under the generalized *logistic* model for some specific choice of parameters. Indeed, consider an instance of the generalized logistic model specified by parameters $\beta_0 = -2.5 \log 7/3, \beta_1 = 5 \log 7/3$ and $\tilde{\nu} = 1$. Then a straightforward verification shows that matrix Π_1^{db} satisfies equation (7a) which specifies the behavior of DB reviewers under the generalized logistic model. Next, observe that for each reviewer $i \in [m]$ and for each paper $j \in \mathcal{J}$ the corresponding entry of the matrix Π_1^{sb} is *larger* than prescribed by the model of SB reviewers under the absence of bias (7b). Similarly, for each paper $j \in \overline{\mathcal{J}}$ the corresponding entry of the matrix Π_1^{sb} is *smaller* than it should be if the bias is absent (7b). Hence, the pair of matrices $\Pi_1^{\text{db}}, \Pi_1^{\text{sb}}$ satisfies the *alternative* hypothesis under the generalized logistic model.

Conversely, consider an instance of the generalized logistic model specified by parameters $\beta_0 = \log 1/3 - 0.625 \log 39/7$, $\beta_1 = 2.5 \log 39/7$ and $\tilde{\nu} = 1.5$. Then a straightforward verification shows that matrix Π_2^{db} satisfies equation (7a) which specifies the behavior of DB reviewer under the generalized logistic model. Recall that at this point we didn't specify how we selected values ν_1, ν_2 in equation (19). In fact, we selected these values such that entries of the matrix Π_2^{sb} satisfy equation (7b) which specifies the behavior of SB reviewers under the generalized logistic model when the bias is absent. Namely, we set

$$\begin{aligned}\nu_1 &= -0.65 + (1 + \exp\{-\beta_0 - \tilde{\nu} - 0.65\beta_1\})^{-1} \\ \nu_2 &= -0.25 + (1 + \exp\{-\beta_0 - \tilde{\nu} - 0.25\beta_1\})^{-1}\end{aligned}$$

As a result, the pair of matrices $\Pi_2^{\text{db}}, \Pi_2^{\text{sb}}$ satisfies the *null* hypothesis under the generalized logistic model.

Finally, the power of the COUNTING test in Figure 5a becomes the Type-I error rate under the instance of the generalized logistic model with $\beta_0 = \log 1/3 - 0.625 \log 39/7$, $\beta_1 = 2.5 \log 39/7$ and $\tilde{\nu} = 1$. Similarly, the Type-I error of the COUNTING test in Figure 5a becomes its power under the instance of the generalized logistic model with $\beta_0 = -2.5 \log 7/3$, $\beta_1 = 5 \log 7/3$ and $\tilde{\nu} = 1$. The same applies to the DISAGREEMENT test and eventually we obtain Figure 5b by simply exchanging the bars in Figure 5a.

E Proofs of main results

In this section we give proofs of our main results.

E.1 Proof of Theorem 1

We proof Theorem 1 in two steps. First, we show the result for the DISAGREEMENT test and then for the COUNTING test. Before we delve into proofs, let us make two observations that we use in this section.

Observations:

- A For every paper $j \in [n]$, if in assignment A^* (Step 2 of Procedure 1) paper j is attributed to reviewers i_1 and i_2 , then the events “reviewer i_1 is allocated to SB and reviewer i_2 is allocated to DB” and “reviewer i_1 is allocated to DB and reviewer i_2 is allocated to SB” are mutually exclusive and happen with probability 0.5 each. This is ensured by Step 3 of Procedure 1 where reviewers are allocated to conditions.
- B By construction of Procedure 1, at least $c \min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ papers each from sets \mathcal{J} and $\overline{\mathcal{J}}$ appear in assignment A^* for some constant c that depends only on parameters λ and μ . Indeed, in cases (a) and (c) of Step 2, all papers are included into assignment A^* and hence our claim holds with $c = 1$. In case (b) $\frac{m}{2} \geq \frac{\lambda}{\mu} n \geq 2 \frac{\lambda}{\mu} \min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ papers are selected and hence our claim holds with $c = \frac{\lambda}{\mu}$.

E.1.1 Proof for DISAGREEMENT test

The proof of Theorem 1 for the DISAGREEMENT test consists of two parts. First, we show that under the null hypothesis defined in Problem 1, for any matrices Π^{db} and $\Pi^{\text{sb}} (= \Pi^{\text{db}})$ and for any assignment A^* constructed by Procedure 1 in Step 2, the test rejects the null with probability at most α . Second, we show that if the number of papers in both \mathcal{J} and $\overline{\mathcal{J}}$ is large enough, then the DISAGREEMENT test satisfies the requirement of non-trivial power.

We prove both parts conditioned on the assignment A^* . The unconditional statement of the theorem then follows from the law of total probability.

Control over Type-I error

Let Π^{db} and $\Pi^{\text{sb}} (= \Pi^{\text{db}})$ be arbitrary matrices that fall under the definition of null hypothesis in Problem 1. Consider arrays U and V constructed in Step 2 of the DISAGREEMENT test from the set of tuples \mathcal{T} passed to the test by Procedure 1. If any of them is empty, the test keeps the null and hence does not commit the Type-I error. Now without loss of generality assume that both U and V are non-empty.

The idea of the proof is to show that under the null hypothesis, entries of arrays U and V are mutually independent and identically distributed. Assume for the moment that it is indeed the case. Then entries of arrays U and V are exchangeable random variables and hence the permutation test with statistic τ defined in Step 3 of Test 1 is guaranteed to provide control over the Type-I error rate for any given significance level $\alpha \in (0, 1)$ and hence the result for Type-I error control follows.

Consider any entry u of array U . Then u is a decision of SB reviewer for some paper $j_t \in \mathcal{J}$, where t is a tuple that corresponds to u . Corresponding SB and DB reviewers disagree in their decisions, that is, $Y_{j_t} \neq X_{j_t}$. Recalling Observation A, we deduce that conditioned on assignment A^* , the symmetry of the null hypothesis guarantees that

$$Y_{j_t} | (Y_{j_t} \neq X_{j_t}) \sim \text{Bernoulli}(0.5). \quad (20)$$

Indeed, given that both Y_{j_t} and X_{j_t} are Bernoulli random variables, one can verify that

$$\mathbb{P}[Y_{j_t} = 1, X_{j_t} = 0] = \mathbb{P}[Y_{j_t} = 0, X_{j_t} = 1],$$

which coupled with the definition of condition probability implies (20).

Hence, entries of array U are Bernoulli random variables with expectation 0.5. Provided that each reviewer contributes at most one decision to \mathcal{T} , entries of U are also independent. The same argument applies to entries of array V and hence we have shown that under the null hypothesis entries of U and V are independent Bernoulli random variables with probability of success 0.5 and thus are exchangeable.

Non-trivial power

Consider any fixed choice of $\delta > 0$ and $\varepsilon > 0$ in the definition of non-trivial power. The goal now is to show that there exists $n_0 = n_0(\varepsilon, \delta)$ such that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then for any matrices Π^{db} and Π^{sb} that satisfy the alternative hypothesis in Problem 1 with margin δ , the DISAGREEMENT test coupled with Procedure 1 is guaranteed to reject the null hypothesis with probability at least $1 - \varepsilon$. Throughout the proof we use c to denote a universal constant and allow its value to change from line to line due to multiplications by some other universal constants. Recall that problem parameters λ, μ and α are treated as constants. For concreteness, throughout the proof we assume that the bias is in favor of papers from \mathcal{J} . The same argument can be repeated in case of bias against papers from \mathcal{J} .

Step 1. Cardinality of U and V .

Let us first show that arrays U and V will with high probability contain order n_0 elements. To this end, recall that for tuple $t \in \mathcal{T}$ we add Y_{j_t} to U if (i) $w_{j_t} = 1$ and (ii) $Y_{j_t} \neq X_{j_t}$. Observation B ensures that \mathcal{T} will contain at least cn_0 tuples that correspond to papers from \mathcal{J} . Consider any such tuple, and let (j_t, i_1, i_2) be a corresponding paper and two reviewers assigned to this paper in assignment A^* . Then conditioned on assignment A^* , $\mathbb{P}[Y_{j_t} \neq X_{j_t}]$ is lower bounded by:

$$\begin{aligned} \mathbb{P}[Y_{j_t} \neq X_{j_t}] &= \frac{1}{2} \left(\pi_{i_1 j_t}^{(\text{sb})} (1 - \pi_{i_2 j_t}^{(\text{db})}) + \pi_{i_2 j_t}^{(\text{db})} (1 - \pi_{i_1 j_t}^{(\text{sb})}) \right) + \frac{1}{2} \left(\pi_{i_2 j_t}^{(\text{sb})} (1 - \pi_{i_1 j_t}^{(\text{db})}) + \pi_{i_1 j_t}^{(\text{db})} (1 - \pi_{i_2 j_t}^{(\text{sb})}) \right) \\ &\geq \frac{1}{2} \left(\pi_{i_1 j_t}^{(\text{sb})} (1 - \pi_{i_2 j_t}^{(\text{db})}) \right) + \frac{1}{2} \left(\pi_{i_2 j_t}^{(\text{sb})} (1 - \pi_{i_1 j_t}^{(\text{db})}) \right) \\ &\stackrel{(i)}{\geq} \frac{1}{2} (\delta^2 + \delta^2) = \delta^2, \end{aligned}$$

where inequality (i) follows from the fact that for any reviewer $i \in [m]$ and for any paper $j \in [n]$ we have $\delta \leq \pi_{ij}^{(\text{sb})} \leq 1$ and $0 \leq \pi_{ij}^{(\text{db})} \leq 1 - \delta$ by the definition of non-trivial power requirement.

The same argument applies to tuples $t \in \mathcal{T}$ that correspond to papers from $\overline{\mathcal{J}}$. Hence, we conclude that for any tuple $t \in \mathcal{T}$ we are guaranteed that $Y_{j_t} \neq X_{j_t}$ with probability at least δ^2 .

Now notice that $|U| = \sum_{t \in \mathcal{T}: w_{j_t} = 1} \mathbb{I}[Y_{j_t} \neq X_{j_t}]$ and hence $\mathbb{E}[|U|] \geq cn_0 \delta^2$. Applying Hoeffding's inequality, we can also derive that for large enough n_0 with probability at least $1 - \frac{\varepsilon}{4}$ we have

$$|U| > cn_0 \delta^2.$$

The same argument applies to V and hence we conclude that with probability at least $1 - \frac{\epsilon}{2}$ we have

$$|U| > cn_0\delta^2 \quad \text{and} \quad |V| > cn_0\delta^2. \quad (21)$$

Step 2. Distribution.

Now we describe the distribution of components of U and V . By construction, the entries of these arrays are independent, so it suffices to study a single component. Consider an entry u of array U and let (j, i_1, i_2) be a corresponding paper and two reviewers assigned to this paper in assignment A^* . For brevity, denote $p = \pi_{i_1 j}^{(\text{sb})} \in (\delta, 1]$, $q = \pi_{i_2 j}^{(\text{db})} \in [0, 1 - \delta)$, $\gamma_1 = p - \pi_{i_1 j}^{(\text{db})}$ and $\gamma_2 = \pi_{i_2 j}^{(\text{sb})} - q$, where $\gamma_1 > \delta$ and $\gamma_2 > \delta$ by definition of non-trivial power requirement. Then, we can derive the following chain of bounds:

$$\begin{aligned} 2\mathbb{P}[u = 1] - 1 &= 2\mathbb{P}[Y_j = 1 | Y_j \neq X_j] - 1 \\ &= \frac{p(1-q)}{p(1-q) + q(1-p)} + \frac{(q + \gamma_2)(1-p + \gamma_1)}{(q + \gamma_2)(1-p + \gamma_1) + (p - \gamma_1)(1-q - \gamma_2)} - 1 \\ &\stackrel{(i)}{\geq} \frac{p(1-q)}{p(1-q) + q(1-p)} + \frac{(q + \delta)(1-p + \delta)}{(q + \delta)(1-p + \delta) + (p - \delta)(1-q - \delta)} - 1 \\ &= \frac{1}{2} \left(\frac{p-q}{p+q-2pq} - \frac{p-q-2\delta}{p+q-2pq+2\delta(\delta+q-p)} \right) \end{aligned}$$

where inequality (i) holds due to monotonicity of the expression over γ_1 and γ_2 and lower bounds $\gamma_1 > \delta$, $\gamma_2 > \delta$.

Optimizing the last expression over $p \in (\delta, 1]$ and $q \in [0, 1 - \delta)$, we obtain

$$2\mathbb{P}[u = 1] - 1 \geq \frac{\delta^2}{\delta^2 + (1 - \delta^2)},$$

and hence $\mathbb{P}[u = 1] \geq \frac{1}{2} + \frac{1}{2} \frac{\delta^2}{\delta^2 + (1 - \delta^2)} = \frac{1}{2} + \gamma$. Similarly, we can show that $\mathbb{P}[v = 1] \leq \frac{1}{2} - \frac{1}{2} \frac{\delta^2}{\delta^2 + (1 - \delta^2)} = \frac{1}{2} - \gamma$, where $\gamma > 0$ is a constant that depends on δ .

Step 3. Permutation.

At this point we are guaranteed that vectors V and U constructed in Step 2 of the DISAGREEMENT test, with probability $1 - \frac{\epsilon}{2}$, contain at least $cn_0\delta^2$ elements and their entries are independent Bernoulli random variables. Moreover, the entries of U have expectations larger than $1/2 + \gamma$ and entries of V have expectations smaller than $1/2 - \gamma$, where γ is independent of n_0 .

Conditioned on $\min\{|V|, |U|\} > cn_0\delta^2$, notice that as n_0 grows, the permutation test for exchangeability of entries of V and U has power growing to 1. Hence, there exists n_0^* such that if $n_0 > n_0^*$, then the permutation test rejects the null with probability at least $1 - \frac{\epsilon}{2}$.

Finally, taking union bound over (i) probability that either of U and V has cardinality smaller than $cn_0\delta^2$ and (ii) probability that the permutation test fails to reject the null given $\min\{|V|, |U|\} > cn_0\delta^2$, we deduce that conditioned on A^* , the requirement of non-trivial power is satisfied. It now remains to notice that the established fact holds for any A^* that is constructed by Procedure 1 and hence Theorem 1(a) holds.

E.1.2 Proof for COUNTING test

Similar to the proof for the DISAGREEMENT test, the proof for the COUNTING test consists of two parts — control over Type-I error and non-trivial power. As in the proof for the DISAGREEMENT test, we prove both parts conditioned on the assignment A^* computed in Step 2 of Procedure 1. The unconditional statement of the theorem then follows from the law of total probability.

Control over Type-I error

Let Π^{db} and $\Pi^{\text{sb}} (= \Pi^{\text{db}})$ be arbitrary matrices that fall under the definition of the null hypothesis in Problem 1. Consider arrays U and V constructed in Step 2 of the COUNTING test. If any of them is empty, the test keeps the null and hence does not commit the Type-I error. Now without loss of generality assume

that both U and V are non-empty. By construction, conditioned on the assignment A^* , entries of arrays U and V are mutually independent and bounded by 1 in absolute value. Moreover, conditioned on A^* the size of arrays U and V is fixed and is not a random variable. Next, we can show that expectation of any entry of arrays U and V is zero. Indeed, consider any arbitrary entry $u \in U$ and let (j, i_1, i_2) be a corresponding paper and reviewers assigned to this paper in assignment A^* . Then:

$$\mathbb{E}[u] = \frac{1}{2} \left(\pi_{i_1 j}^{(\text{sb})} - \pi_{i_2 j}^{(\text{db})} \right) + \frac{1}{2} \left(\pi_{i_2 j}^{(\text{sb})} - \pi_{i_1 j}^{(\text{db})} \right) = 0,$$

where two terms correspond to two equiprobable allocations of reviewers i_1 and i_2 to conditions and the last equality follows from the fact that under the null hypothesis $\Pi^{\text{sb}} = \Pi^{\text{db}}$. Hence, we conclude that the expectation of test statistic γ equals 0. Independence and boundedness of entries of arrays V and U ensure that the test statistic γ is sub-Gaussian random variable with noise parameter σ given by

$$\sigma^2 = |U|^{-1} + |V|^{-1}.$$

Finally, applying Hoeffding's inequality we deduce that

$$\mathbb{P} \left[|\gamma| > \sqrt{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha} \right] \leq 2 \exp \left\{ -\frac{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha}{2(|U|^{-1} + |V|^{-1})} \right\} = \alpha,$$

which concludes the proof.

Non-trivial power

Consider any fixed choice of $\delta > 0$ and $\varepsilon > 0$ in the definition of non-trivial power. The goal now is to show that there exists $n_0 > 0$ such that if $\min\{|\mathcal{J}|, |\bar{\mathcal{J}}|\} > n_0$, then for any matrices Π^{db} and Π^{sb} that satisfy the alternative hypothesis in Problem 1 with margin δ , the COUNTING test coupled with Procedure 1 rejects the null hypothesis with probability at least $1 - \varepsilon$. Throughout the proof we use c to denote a universal constant and allow its value to change from line to line due to multiplications by some other universal constants. Recall that problem parameters λ, μ and α are treated as constants. For concreteness, suppose that there is a bias in favor of papers that satisfy the property of interest.

We now consider an arbitrary instance of the bias testing problem with matrices Π^{sb} and Π^{db} that fall under the definition of non-trivial power. First, Observation B ensures that the set \mathcal{T} passed to the COUNTING algorithm is such that the resulting vectors U and V contain at least cn_0 elements each. Next, let $\gamma_1 = \frac{1}{|U|} \sum_{u \in U} u$ and $\gamma_2 = \frac{1}{|V|} \sum_{v \in V} v$, in this notation the test statistic is defined as $\gamma = \gamma_1 - \gamma_2$. Conditioned on the assignment A^* , we have:

$$\mathbb{E}[\gamma_1] = \frac{1}{|U|} \sum_{u \in U} \mathbb{E}[u] \geq \delta.$$

Indeed, for any arbitrary entry u of array U let (j, i_1, i_2) be corresponding paper and reviewers assigned to this paper in assignment A^* . Then requirement of non-trivial power guarantees that

$$\mathbb{E}[u] = \frac{1}{2} \left(\pi_{i_1 j}^{(\text{sb})} - \pi_{i_2 j}^{(\text{db})} \right) + \frac{1}{2} \left(\pi_{i_2 j}^{(\text{sb})} - \pi_{i_1 j}^{(\text{db})} \right) \geq \delta + \frac{1}{2} \left(\pi_{i_1 j}^{(\text{db})} - \pi_{i_2 j}^{(\text{db})} \right) + \frac{1}{2} \left(\pi_{i_2 j}^{(\text{db})} - \pi_{i_1 j}^{(\text{db})} \right) = \delta.$$

Similarly,

$$\mathbb{E}[\gamma_2] = \frac{1}{|V|} \sum_{v \in V} \mathbb{E}[v] \leq -\delta.$$

Applying Hoeffding's inequality we obtain:

$$\mathbb{P}[\gamma_1 - \gamma_2 < \delta] \leq \mathbb{P}[\gamma_1 - \gamma_2 < \mathbb{E}[\gamma_1 - \gamma_2] - \delta] \leq \exp \left(-\frac{\delta^2}{2(|V|^{-1} + |U|^{-1})} \right) \leq \exp(-c\delta^2 n_0).$$

On the other hand, the threshold for rejecting the null is such that

$$\sqrt{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha} \leq c\sqrt{\frac{1}{n_0}}.$$

Finally, setting $n_0 = c\frac{\log 1/\varepsilon}{\delta^2}$, we ensure that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then the COUNTING algorithm with probability at least $1 - \varepsilon$ rejects the null for any matrices $\Pi^{\text{sb}}, \Pi^{\text{db}}$ that satisfy alternative hypothesis with margin δ .

E.2 Proof of Theorem 2

We proof Theorem 2 separately for the DISAGREEMENT test and for the COUNTING tests.

E.2.1 Proof for DISAGREEMENT test

Again, the proof is presented in two parts: control over Type-I error and non-trivial power. The conceptual difference from the proof of the corresponding result for absolute bias problem is that now the parametric relationships (7a) and (7b) allow us to avoid conditioning on the assignment A^* .

Control over Type-I error

Let Π^{db} and Π^{sb} be arbitrary matrices generated from the generalized logistic model under the absence of bias. Consider arrays U and V constructed in Step 2 of the DISAGREEMENT test from the set of tuples \mathcal{T} passed to the test by Procedure 1. If any of them is empty, the test keeps the null and hence does not commit the Type-I error. Now without loss of generality assume that both arrays U and V are non-empty. Following the idea of the proof of Theorem 1, we need to show that entries of arrays U and V are exchangeable random variables. First, the mutual independence follows from construction of the set \mathcal{T} . Second, using equations (7a) and (7b), we deduce that for any paper $j \in [n]$ and for any reviewer $i \in [m]$:

$$\log \frac{\pi_{ij}^{(\text{sb})}(1 - \pi_{ij}^{(\text{db})})}{\pi_{ij}^{(\text{db})}(1 - \pi_{ij}^{(\text{sb})})} = \tilde{\nu}.$$

Noticing that $\pi_{ij}^{(\text{sb})}$ and $\pi_{ij}^{(\text{db})}$ under the generalized logistic model are independent of reviewer's identity, we drop index i from the above equation. Now we consider any entry u of array U together with a corresponding tuple $t = (j_t, Y_{j_t}, X_{j_t}, w_{j_t})$ and conclude that:

$$\begin{aligned} \mathbb{P}[u = 1] &= \mathbb{P}[Y_{j_t} = 1 | Y_{j_t} \neq X_{j_t}] \\ &= \frac{\pi_{j_t}^{(\text{sb})}(1 - \pi_{j_t}^{(\text{db})})}{\pi_{j_t}^{(\text{sb})}(1 - \pi_{j_t}^{(\text{db})}) + \pi_{j_t}^{(\text{db})}(1 - \pi_{j_t}^{(\text{sb})})} \\ &= \frac{1}{1 + \frac{\pi_{j_t}^{(\text{db})}(1 - \pi_{j_t}^{(\text{sb})})}{\pi_{j_t}^{(\text{sb})}(1 - \pi_{j_t}^{(\text{db})})}} \\ &= \frac{1}{1 + e^{-\tilde{\nu}}}. \end{aligned} \tag{22}$$

Importantly, the value of the paper representation q_j does not appear in equation (22), implying that entries of array U are identically distributed. Applying the same argument to entries of array V we deduce that entries of arrays U and V are exchangeable random variables and hence the permutation test with the test statistic τ defined in Step 3 of Test 1 is guaranteed to control for the Type-I error rate at any given significance level $\alpha \in (0, 1)$ which concludes the proof.

Non-trivial power

Consider any fixed choice of $\delta > 0$ and $\varepsilon > 0$ in the definition of non-trivial power. The goal now is to show that there exists $n_0 = n_0(\varepsilon, \delta)$ such that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then for any matrices Π^{db} and Π^{sb} generated

from the generalized logistic model that satisfy the alternative hypothesis in Problem 2 with margin δ , the DISAGREEMENT test coupled with Procedure 1 is guaranteed to reject the null hypothesis with probability at least $1 - \varepsilon$. Throughout the proof we use c to denote a universal constant and allow its value to change from line to line due to multiplications by some other universal constants. Recall that problem parameters λ, μ and α are treated as constants. For concreteness, throughout the proof we assume that the bias is in favor of papers from \mathcal{J} . The same argument can be repeated in case of bias against papers from \mathcal{J} .

Step 1. Cardinality of U and V .

Consider any matrices Π^{sb} and Π^{db} generated from logistic regression model that satisfy the alternative hypothesis in Problem 2 with margin δ . First, we notice that scores $q_j, j \in [n]$, and coefficients β_0, β_1 are bounded in absolute value by some constant $\tilde{\Delta}$, and hence using equation (7a) we conclude that for all $(i, j) \in [n] \times [m]$

$$\pi_{ij}^{(\text{db})} \in (\ell, b) \quad \forall j \in [n], \quad (23)$$

where $0 < \ell < b < 1$ and values of ℓ and b are determined by $\tilde{\Delta}$. Now consider any tuple $t = (j_t, Y_{i_1 j_t}, X_{i_2 j_t}, w_{j_t})$ from the set of tuples \mathcal{T} . Then

$$\begin{aligned} \mathbb{P}[Y_{j_t} \neq X_{j_t}] &= \pi_{i_1 j_t}^{(\text{sb})}(1 - \pi_{i_2 j_t}^{(\text{db})}) + \pi_{i_2 j_t}^{(\text{db})}(1 - \pi_{i_1 j_t}^{(\text{sb})}) \\ &\geq \min\{\pi_{i_2 j_t}^{(\text{db})}, 1 - \pi_{i_2 j_t}^{(\text{db})}\} \left(\pi_{i_1 j_t}^{(\text{sb})} + 1 - \pi_{i_1 j_t}^{(\text{sb})} \right) \\ &= \min\{\pi_{i_2 j_t}^{(\text{db})}, 1 - \pi_{i_2 j_t}^{(\text{db})}\} \\ &\geq \min\{\ell, 1 - b\}, \end{aligned}$$

where the last inequality follows from equation (23). Applying Hoeffding's inequality in the same way as we did in the proof of Theorem 1 to get the bound (21), we deduce that with probability at least $1 - \frac{\varepsilon}{2}$, cardinalities of arrays U and V are at least cn_0 for some constant c that may depend on δ and $\tilde{\Delta}$.

Step 2. Distribution

By definition of non-trivial power requirement, it must be the case that for all $(i, j) \in [m] \times [n]$ we have $|\pi_{ij}^{(\text{sb})} - f_0(\pi_{ij}^{(\text{db})})| > \delta$, where function f_0 belongs to class $\tilde{\mathcal{F}}_{\tilde{\Delta}}$ defined in (8) and for all $(i, j) \in [m] \times [n]$ satisfies:

$$\log \frac{f_0(\pi_{ij}^{(\text{db})})}{1 - f_0(\pi_{ij}^{(\text{db})})} = \beta_0 + \tilde{\nu} + \beta_1 q_j \quad (24a)$$

$$= \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} + \tilde{\nu}, \quad (24b)$$

for some value of $\tilde{\nu} \in (-\tilde{\Delta}, \tilde{\Delta})$. Observe that values $\beta_0, \tilde{\nu}, \beta_1, q_j$ in the RHS of equation (24a) are bounded in absolute value by constant $\tilde{\Delta}$. Next, recall that the definition of the non-trivial power requirement ensures that for each reviewer $i \in [m]$ it must be the case that (a) for each paper $j \in \mathcal{J}$ we have $f_0(\pi_{ij}^{(\text{db})}) < 1 - \delta$ and (b) for each paper $j \in \bar{\mathcal{J}}$ we have $f_0(\pi_{ij}^{(\text{db})}) > \delta$. Finally, we are guaranteed that for any pair of reviewer $i \in [m]$ and paper $j \in [n]$ we have

$$f_0(\pi_{ij}^{(\text{db})}) \in \begin{cases} (\ell', \min\{b', 1 - \delta\}) & \text{if } j \in \mathcal{J} \\ (\max\{\ell', \delta\}, b') & \text{if } j \in \bar{\mathcal{J}}. \end{cases}$$

Notice that constants ℓ' and b' are such that $0 < \ell' < b' < 1$ and may be different from ℓ and b , because in equation (24a) we have additional term $\tilde{\nu}$ which is absent in (7a).

Let us now define two quantities d_1 and d_2 as

$$d_1 = \inf_{t \in (\ell', \min\{b', 1-\delta\})} \left(\log \frac{t + \delta}{1 - (t + \delta)} - \log \frac{t}{1 - t} \right) \quad (25a)$$

$$d_2 = \inf_{t \in (\max\{\ell', \delta\}, b')} \left(\log \frac{t}{1 - t} - \log \frac{t - \delta}{1 - (t - \delta)} \right). \quad (25b)$$

Notice that both quantities d_1 and d_2 are some functions of δ and $\tilde{\Delta}$ and are strictly positive, because function $\log \frac{x}{1-x}$ is strictly increasing on the interval $(0, 1)$ with its derivative being lower bounded by $c > 0$, where c is independent of problem parameters.

Putting together equations (24a) - (25b), we now show that for each reviewer $i \in [m]$ the definition of non-trivial power requirement ensures that for each paper $j \in \mathcal{J}$

$$\log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} \geq \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} + \tilde{\nu} + d_1, \quad (26)$$

and for each paper $j \in \overline{\mathcal{J}}$

$$\log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} \leq \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} + \tilde{\nu} - d_2. \quad (27)$$

Consider any arbitrary entry u of array U and the corresponding tuple $(j_t, Y_{i_1 j_t}, X_{i_2 j_t}, w_{j_t})$. Then,

$$\begin{aligned} \mathbb{P}[u = 1] &= \mathbb{P}[Y_{i_1 j_t} = 1 | Y_{i_1 j_t} \neq X_{i_2 j_t}] \\ &= \frac{\pi_{i_1 j_t}^{(\text{sb})} (1 - \pi_{i_2 j_t}^{(\text{db})})}{\pi_{i_1 j_t}^{(\text{sb})} (1 - \pi_{i_2 j_t}^{(\text{db})}) + \pi_{i_2 j_t}^{(\text{db})} (1 - \pi_{i_1 j_t}^{(\text{sb})})} \\ &= \frac{1}{1 + \frac{\pi_{i_2 j_t}^{(\text{db})} (1 - \pi_{i_1 j_t}^{(\text{sb})})}{\pi_{i_1 j_t}^{(\text{sb})} (1 - \pi_{i_2 j_t}^{(\text{db})})}} \\ &\geq \frac{1}{1 + e^{-\tilde{\nu} - d_1}}, \end{aligned}$$

where the last inequality follows from (26). Similarly, using (27) we show that for each entry v of array V

$$\mathbb{P}[v = 1] \leq \frac{1}{1 + e^{-\tilde{\nu} + d_2}}.$$

Step 3. Permutation.

At this point we are guaranteed that vectors V and U constructed in Step 2 of the DISAGREEMENT test, with probability $1 - \frac{\varepsilon}{2}$, contain at least cn_0 elements and their entries are independent Bernoulli random variables. Moreover, the entries of U have expectations larger than $\frac{1}{1+e^{-\tilde{\nu}}} + \gamma$ and entries of V have expectations smaller than $\frac{1}{1+e^{-\tilde{\nu}}} - \gamma$, where γ is independent of n_0 , but depends on δ and $\tilde{\Delta}$.

Conditioned on $\min\{|V|, |U|\} > cn_0$, notice that as n_0 grows, the permutation test for exchangeability of entries of V and U has power growing to 1. Hence, there exists n_0^* such that if $n_0 > n_0^*$, then the permutation test rejects the null with probability at least $1 - \frac{\varepsilon}{2}$.

Finally, taking union bound over (i) probability that either of U and V has cardinality smaller than cn_0 and (ii) probability that the permutation test fails to reject the null given $\min\{|V|, |U|\} > cn_0$, we deduce that the requirement of non-trivial power is satisfied.

E.2.2 Proof for COUNTING test

We give a proof for an extended version of the generalized linear model in which for each $(i, j) \in [n] \times [m]$ we substitute q_j with q_{ij} , thus allowing subjectivity of reviewers. In the proof we will be using two observations we made in the beginning of Appendix E.1. As in the proof of Theorem 1, we prove the result conditioned on the assignment A^* constructed in Step 2 of Procedure 1. The unconditional statement of the theorem then follows from the law of total probability.

Control over Type-I error

Let Π^{db} and Π^{sb} be arbitrary matrices generated under the generalized linear model that fall under the null hypothesis in Problem 2. Consider arrays U and V constructed in Step 2 of the COUNTING test. If any of them is empty, the test keeps the null and hence does not commit the Type-I error. Now without loss of generality assume that both U and V are non-empty. By construction, conditioned on the assignment A^* , entries of arrays U and V are mutually independent and bounded by 1 in absolute value. Moreover, conditioned on A^* the size of arrays U and V is fixed and is not a random variable. Next, for any arbitrary entry $u \in U$ let (j, i_1, i_2) be a corresponding paper and reviewers assigned to this paper in assignment A^* . Then,

$$\begin{aligned} \mathbb{E}[u] &= \frac{1}{2} \left(\pi_{i_1 j}^{(\text{sb})} - \pi_{i_2 j}^{(\text{db})} \right) + \frac{1}{2} \left(\pi_{i_2 j}^{(\text{sb})} - \pi_{i_1 j}^{(\text{db})} \right) \\ &= \frac{1}{2} (q_{i_1 j} + \nu - q_{i_2 j}) + \frac{1}{2} (q_{i_2 j} + \nu - q_{i_1 j}) \\ &= \nu. \end{aligned}$$

Similarly, it follows that for any arbitrary entry $v \in V$:

$$\mathbb{E}[v] = \nu.$$

Hence, we conclude that the expectation of the test statistic γ equals 0. Independence and boundedness of entries of arrays V and U ensure that the test statistic γ is sub-Gaussian random variable with noise parameter σ given by

$$\sigma^2 = |U|^{-1} + |V|^{-1}.$$

Finally, applying Hoeffding's inequality we deduce that

$$\mathbb{P} \left[|\gamma| > \sqrt{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha} \right] \leq 2 \exp \left\{ -\frac{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha}{2(|U|^{-1} + |V|^{-1})} \right\} = \alpha,$$

which concludes the proof.

Non-trivial power

Consider any fixed choice of $\delta > 0$ and $\varepsilon > 0$ in the definition of non-trivial power. The goal now is to show that there exists $n_0 > 0$ such that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then for any matrices Π^{db} and Π^{sb} generated under the generalized linear model that satisfy the alternative hypothesis in Problem 2 with margin δ , the COUNTING test coupled with Procedure 1 rejects the null hypothesis with probability at least $1 - \varepsilon$. Throughout the proof we use c to denote a universal constant and allow its value to change from line to line due to multiplications by some other universal constants. Recall that problem parameters λ, μ and α are treated as constants. For concreteness, suppose that there is a bias in favor of papers that satisfy the property of interest.

We now consider an arbitrary instance of the bias testing problem with matrices Π^{sb} and Π^{db} that fall under the definition of non-trivial power. First, Observation B ensures that the set \mathcal{T} passed to the COUNTING algorithm is such that resulting vectors U and V contain at least cn_0 elements each. Next, let $\gamma_1 = \frac{1}{|U|} \sum_{u \in U} u$ and $\gamma_2 = \frac{1}{|V|} \sum_{v \in V} v$, in this notation the test statistic is defined as $\gamma = \gamma_1 - \gamma_2$. Conditioned on the assignment

A^* , we have:

$$\mathbb{E}[\gamma_1] = \frac{1}{|U|} \sum_{u \in U} \mathbb{E}[u] \geq \nu + \delta.$$

Indeed, for any arbitrary entry u of array U let (j, i_1, i_2) be corresponding paper and reviewers assigned to this paper in assignment A^* . Then the definition of the non-trivial power guarantees that

$$\begin{aligned} \mathbb{E}[u] &= \frac{1}{2} \left(\pi_{i_1 j}^{(\text{sb})} - \pi_{i_2 j}^{(\text{db})} \right) + \frac{1}{2} \left(\pi_{i_2 j}^{(\text{sb})} - \pi_{i_1 j}^{(\text{db})} \right) \\ &\geq \frac{1}{2} (q_{i_1 j} + \nu + \delta - q_{i_2 j}) + \frac{1}{2} (q_{i_2 j} + \nu + \delta - q_{i_1 j}) \\ &= \nu + \delta. \end{aligned}$$

Similarly,

$$\mathbb{E}[\gamma_2] = \frac{1}{|V|} \sum_{v \in V} \mathbb{E}[v] \leq \nu - \delta.$$

Applying Hoeffding's inequality we obtain:

$$\mathbb{P}[\gamma_1 - \gamma_2 < \delta] \leq \mathbb{P}[\gamma_1 - \gamma_2 < \mathbb{E}[\gamma_1 - \gamma_2] - \delta] \leq \exp\left(-\frac{\delta^2}{2(|V|^{-1} + |U|^{-1})}\right) \leq \exp(-c\delta^2 n_0).$$

On the other hand, the threshold for acceptance is such that

$$\sqrt{2(|U|^{-1} + |V|^{-1}) \log 2/\alpha} \leq c\sqrt{\frac{1}{n_0}}.$$

Finally, setting $n_0 = c\frac{\log 1/\varepsilon}{\delta^2}$, we ensure that if $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\} > n_0$, then the COUNTING algorithm with probability at least $1 - \varepsilon$ rejects the null for any matrices $\Pi^{\text{sb}}, \Pi^{\text{db}}$ that satisfy the alternative hypothesis with margin δ .

E.3 Proof of Theorem 3

Assume that the premises of Theorem 3 are satisfied, that is, there exist functions $g, h \in \mathcal{F}$ and values $0 \leq x_1 < x_2 \leq 1$ such that $g(x_1) < h(x_1)$ and $g(x_2) > h(x_2)$.

The high-level idea of the proof is to construct matrices Π^{db} and Π^{sb} which simultaneously satisfy the null hypothesis of Problem 2 specified by some function $f_0 \in \mathcal{F}$ and the alternative hypothesis of Problem 2 specified by another function $f'_0 \in \mathcal{F}$ with margin $\delta > 0$. If such matrices exist, then there exist two instances of a bias testing problem — one with presence of bias and the other with absence of bias — such that the distributions of the reviewers' decisions for these two instances coincide. Hence, any test that uniformly controls for the Type-I error rate at the level α for every $f_0 \in \mathcal{F}$ must under the second instance have power upper bounded by α and thus violate the requirement of non-trivial power over the class of functions \mathcal{F} .

We begin with building a matrix Π^{db} . For any reviewer $i \in [m]$ and for any paper $j \in [n]$ we let

$$\pi_{ij}^{(\text{db})} = \begin{cases} x_1 & \text{if } w_j = 1 \\ x_2 & \text{if } w_j = -1. \end{cases}$$

Next, we define Π^{sb} as follows. For any reviewer $i \in [m]$ and for any paper $j \in [n]$

$$\pi_{ij}^{(\text{sb})} = h(\pi_{ij}^{(\text{db})}).$$

By construction matrices Π^{sb} and Π^{db} satisfy the null hypothesis specified by function $h \in \mathcal{F}$. On the other hand, notice that for each paper $j \in \mathcal{J}$ we have

$$\pi_{ij}^{(\text{sb})} = h(x_1) > g(x_1) = g(\pi_{ij}^{(\text{db})}),$$

and for each paper $j \in \bar{\mathcal{J}}$ we have

$$\pi_{ij}^{(\text{sb})} = h(x_2) < g(x_2) = g(\pi_{ij}^{(\text{db})}).$$

Hence, matrices Π^{db} and Π^{sb} also satisfy the alternative hypothesis specified by function g . Moreover, Π^{db} and Π^{sb} satisfy this alternative with margin $\delta = \min\{|h(x_1) - g(x_1)|, |h(x_2) - g(x_2)|\} > 0$. We now conclude the proof by noting that our construction holds for any choice of parameters λ, μ, n, m and hence the requirement of non-trivial power must be violated by any testing algorithm that controls for Type-I error at the level $\alpha \in (0, 1)$.

E.4 Proof of Corollary 1

To prove Corollary 1, we consider any choice of parameters $\Delta \in (0, 0.5)$ and $\tilde{\Delta} > 0$ and construct two functions f_0 and \tilde{f}_0 together with two numbers $0 \leq x_1 < x_2 \leq 1$ such that

- (i) Functions f_0 and \tilde{f}_0 describe the behavior of reviewers under the absence of bias under the generalized linear and generalized logistic models respectively, that is, $f_0 \in \mathcal{F}_\Delta$, $\tilde{f}_0 \in \tilde{\mathcal{F}}_{\tilde{\Delta}}$, where class \mathcal{F}_Δ is defined by equation (5) and class $\tilde{\mathcal{F}}_{\tilde{\Delta}}$ is specified in equation (8)
- (ii) Values x_1 and x_2 are such that:
 - (a) One can select parameters $q_j^*, j \in [n]$, that fall under the definition of the generalized linear model such that matrix Π^{db} generated according to the equation (4a) satisfies the following equation:

$$\pi_{ij}^{(\text{db})} = \begin{cases} x_1 & \text{if } j \in \mathcal{J} \\ x_2 & \text{if } j \in \bar{\mathcal{J}}. \end{cases} \quad (28)$$

- (b) One can select parameters $q_j^*, j \in [n]$, and β_0, β_1 that fall under the definition of the generalized logistic model such that matrix Π^{db} generated according to the equation (7a) satisfies the equation (28).
- (iii) Functions f_0 and \tilde{f}_0 are such that

$$\text{sign}\left(f_0(x_1) - \tilde{f}_0(x_1)\right) \times \text{sign}\left(f_0(x_2) - \tilde{f}_0(x_2)\right) = -1,$$

where $\text{sign}(\cdot)$ is the sign function. That is, at x_1 the function f_0 is strictly larger than \tilde{f}_0 and at x_2 the function f_0 is strictly smaller than \tilde{f}_0 , or vice versa.

Assume for the moment that conditions (i)-(iii) are satisfied and consider the matrix Π^{db} whose entries are given by equation (28). Then one can select values of papers' representations $q_j, j \in [n]$, such that Π^{db} satisfies the model of DB reviewers in the generalized linear model (4a). Similarly, there exists another choice of papers' representations $q'_j, j \in [n]$, and parameters β_0, β_1 , such that the same matrix Π^{db} satisfies the model of DB reviewers in the generalized logistic model (7a). Now define matrix Π_1^{sb} whose entries for each $(i, j) \in [m] \times [n]$ are given by:

$$\pi_{ij}^{(\text{sb})} = f_0(\pi_{ij}^{(\text{db})})$$

and matrix Π_2^{sb} whose entries for each $(i, j) \in [m] \times [n]$ are given by:

$$\pi_{ij}^{(\text{sb})} = \tilde{f}_0(\pi_{ij}^{(\text{db})}).$$

Matrices $\Pi^{\text{db}}, \Pi_1^{\text{sb}}$ satisfy the null hypothesis under the generalized *linear* model specified by the function f_0 . Moreover, condition (iii) ensures that they *simultaneously* satisfy the alternative hypothesis under the generalized *logistic* model specified by the function \tilde{f}_0 with margin $\delta = \min\{|f_0(x_1) - \tilde{f}_0(x_1)|, |f_0(x_2) - \tilde{f}_0(x_2)|\}$. Hence, if the testing procedure ψ_2 (which has a non-trivial power under the generalized logistic model) is given decisions of SB and DB reviewers sampled according to the pair of matrices $\Pi^{\text{db}}, \Pi_1^{\text{sb}}$, then it will reject the null hypothesis with probability that goes to 1 as the minimum of $|\mathcal{J}|$ and $|\tilde{\mathcal{J}}|$ grows. Finally, given that matrices Π^{db} and Π^{sb} solely determine the distribution of observed reviewers' decisions, our construction implies that under the generalized linear model procedure ψ_2 does not control for the Type-I error rate at any level $\alpha < 1$.

A similar argument applies to the pair of matrices $\Pi^{\text{db}}, \Pi_2^{\text{sb}}$, and it follows that under the generalized logistic model the procedure ψ_1 does not control for the Type-I error rate at any level $\alpha < 1$.

To conclude the proof it remains to find $x_1, x_2, f_0, \tilde{f}_0$ that satisfy aforementioned conditions (i)-(iii). To this end, let us define quantities γ_1, γ_2 :

$$\begin{aligned} \gamma_1 &= \max \left\{ \Delta, \left(1 + \exp\{\tilde{\Delta} + \tilde{\Delta}^2\} \right)^{-1} \right\} \\ \gamma_2 &= \min \left\{ 1 - \Delta, \left(1 + \exp\{-\tilde{\Delta} - \tilde{\Delta}^2\} \right)^{-1} \right\}. \end{aligned}$$

Notice that the value of γ_1 is by definition smaller than 0.5. Moreover, for each pair $(i, j) \in [m] \times [n]$ it gives a lower bound on the value $\pi_{ij}^{(\text{db})}$ that can be generated from both the generalized linear (with parameter Δ) and generalized logistic (with parameter $\tilde{\Delta}$) models. Likewise, the value of γ_2 is at least 0.5 and gives the corresponding upper bound. Hence, any values of x_1, x_2 such that $\gamma_1 < x_1 < x_2 < \gamma_2$ satisfy the condition (ii).

Next, find values $\nu \in (0, \Delta)$ and $\tilde{\nu} \in (0, \tilde{\Delta})$ such that for functions $h_\nu \in \mathcal{F}_\Delta$ and $g_{\tilde{\nu}} \in \tilde{\mathcal{F}}_{\tilde{\Delta}}$ defined in equations (5) and (8) respectively the following equality holds:

$$h_\nu(0.5) = g_{\tilde{\nu}}(0.5).$$

Observe that such values must exist because h_ν and $g_{\tilde{\nu}}$ are continuous functions of ν and $\tilde{\nu}$ respectively and

$$\lim_{\nu \rightarrow +0} h_\nu(0.5) = \lim_{\tilde{\nu} \rightarrow +0} g_{\tilde{\nu}}(0.5) = 0.5.$$

Consider now two possible cases:

Case 1. Functions h_ν and $g_{\tilde{\nu}}$ are such that there exist two points $y \in (\gamma_1, 0.5)$ and $z \in (0.5, \gamma_2)$ for which the following equation holds:

$$\text{sign}\left(h_\nu(y) - g_{\tilde{\nu}}(y)\right) \times \text{sign}\left(h_\nu(z) - g_{\tilde{\nu}}(z)\right) = -1, \quad (29)$$

Observe that in this case conditions (i)-(iii) are satisfied by the choice $f_0 = h_\nu, \tilde{f}_0 = g_{\tilde{\nu}}, x_1 = y, x_2 = z$ and hence the result of the theorem follows.

Case 2. Functions h_ν and $g_{\tilde{\nu}}$ are such that h_ν is a tangent line to $g_{\tilde{\nu}}$ at 0.5. This case reduces to the Case 1 by setting $\nu' = \nu - \varepsilon$ for a sufficiently small $\varepsilon \in (0, \nu)$. Indeed, if ε is sufficiently small, then due to strict concavity and differentiability of the function $g_{\tilde{\nu}}$, by shifting the tangent line down we ensure that there exist points $y = 0.5$ and $z \in (0.5, \gamma_2)$ such that

$$\text{sign}\left(h_{\nu'}(y) - g_{\tilde{\nu}}(y)\right) \times \text{sign}\left(h_{\nu'}(z) - g_{\tilde{\nu}}(z)\right) = -1.$$

Hence, we can satisfy conditions (i)-(iii) by setting $f_0 = h_{\nu'}, \tilde{f}_0 = g_{\tilde{\nu}}, x_1 = 0.5, x_2 = z$.

To conclude the proof, we notice that Cases 1 and 2 are complementary, because function $g_{\tilde{\nu}}$ is differentiable and strictly concave on the interval $(0, 1)$ and function h_ν is a linear function.

E.5 Proofs of auxiliary results

In this section we give proofs for auxiliary results stated in appendix.

E.5.1 Proof of Proposition 1

We prove Lemma 1 by straightforward verification. First, let Y_{ij} be generated from model (11b) and $X_{i'j}$ be generated from model (11a). Then

$$\mathbb{E}[Y_{ij} - X_{i'j}] = q_j + \beta_0^{(\text{sb})} + \sum_{\ell \in [k]} \beta_\ell^{(\text{sb})} w_j^{(\ell)} - q_j = \beta_0^{(\text{sb})} + \sum_{\ell \in [k]} \beta_\ell^{(\text{sb})} w_j^{(\ell)}.$$

Similarly, let Y_{ij} be generated from model (12b) and $X_{i'j}$ be generated from model (12a). Then

$$\begin{aligned} \mathbb{E}[Y_{ij} | Y_{ij} \neq X_{i'j}] &= \frac{\pi_{ij}^{(\text{sb})}(1 - \pi_{i'j}^{(\text{db})})}{\pi_{ij}^{(\text{sb})}(1 - \pi_{i'j}^{(\text{db})}) + \pi_{i'j}^{(\text{db})}(1 - \pi_{ij}^{(\text{sb})})} \\ &= \left[1 + \frac{\pi_{i'j}^{(\text{db})}(1 - \pi_{ij}^{(\text{sb})})}{\pi_{ij}^{(\text{sb})}(1 - \pi_{i'j}^{(\text{db})})} \right]^{-1} \\ &= \left[1 + \exp \left\{ - \left(\beta_0^{(\text{sb})} + \sum_{\ell \in [k]} \beta_{\ell+1}^{(\text{sb})} w_j^{(\ell)} - \beta_0^{(\text{db})} \right) \right\} \right]^{-1}, \end{aligned}$$

and hence

$$\log \frac{\mathbb{E}[Y_{ij} | (Y_{ij} \neq X_{i'j})]}{1 - \mathbb{E}[Y_{ij} | (Y_{ij} \neq X_{i'j})]} = \beta_0^{(\text{sb})} - \beta_0^{(\text{db})} + \sum_{\ell \in [k]} \beta_{\ell+1}^{(\text{sb})} w_j^{(\ell)}.$$

E.5.2 Proof of Lemma 1

Consider any assignment of papers to SB reviewers that satisfy (λ, μ) -constraint with $\lambda > \mu$. Then pick any subset of papers $\mathcal{P} \subseteq [n]$ and denote a set of SB reviewers who are assigned to at least one paper from \mathcal{P} as \mathcal{R}_{SB} . Then one can notice that

$$|\mathcal{R}_{\text{SB}}| \geq \frac{\lambda |\mathcal{P}|}{\mu} \geq |\mathcal{P}|,$$

and hence by Hall's theorem there exists a matching that maps each paper to one reviewer such that each reviewer is matched to at most one paper. This matching is computed in Step 2 of Algorithm 1.

The same argument applies to DB reviewers and hence, joining these two matchings, the algorithm in Step 4 constructs a set of tuples \mathcal{T} where for each paper $j \in [n]$ there exists a tuple that corresponds to this paper.

E.5.3 Proof of Lemma 2

Consider any assignments of papers to SB and DB reviewers that satisfy (λ, μ) -constraints. Let γ be a maximum integer that satisfies inequality

$$\gamma \leq \min \left\{ \frac{|\mathcal{J}|}{4\mu}, \frac{|\overline{\mathcal{J}}|}{4\mu} \right\}.$$

Without loss of generality, assume that $\gamma > 1$. Given that μ and λ are treated as constants and that we only need to prove the result for large enough $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$, we ignore the cases when $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ is small.

Consider a graph G before the first iteration of Steps 2 - 4 of Algorithm 2. Each paper in this graph is connected to λ SB and λ DB reviewers such that each reviewer is connected to at most μ papers.

Now let (i_1, j, i_2) and (i'_1, j', i'_2) be triples found in the first iteration of the algorithm. These triples exists provided that $\gamma > 1$. Then in Step 4 we remove reviewers i_1, i'_1, i_2, i'_2 and corresponding edges from graph G . One can see that these reviewers are connected to at most 4μ papers in total and hence before the second iteration of Steps 2 - 4 graph G will have at least $|\mathcal{J}| - 4\mu \geq 4\mu(\gamma - 1)$ papers from \mathcal{J} and $|\overline{\mathcal{J}}| - 4\mu \geq 4\mu(\gamma - 1)$ papers from $\overline{\mathcal{J}}$ that are connected to λ SB and λ DB remaining reviewers and each of the remaining reviewers (there must be at least $8\lambda(\gamma - 1)$ SB and $8\lambda(\gamma - 1)$ DB reviewers) will be connected to at most μ papers.

By induction we can show that in the first γ iterations of Steps 2 - 4 the greedy algorithm will be able to find non-empty triples in Steps 2 and 3. Hence the resulting set of tuples \mathcal{T} will contain at least γ tuples that correspond to papers from \mathcal{J} and at least γ tuples that correspond to papers from $\overline{\mathcal{J}}$. We then conclude the proof noticing that $\gamma = c \min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$, where c is a constant that depends only on μ .

E.5.4 Proof of Proposition 2

The proof of Proposition 2 follows the idea of the proof of Theorem 1 with some changes which we now discuss. Consider any set of triples \mathcal{C} such that (i) each triple $c \in \mathcal{C}$ is of the form (j, i_1, i_2) (one paper and two reviewers) and (ii) each reviewer $i \in [m]$ appears in at most one triple. Let \mathbb{C} denote a collection of all such sets of triples. Then any set of tuples \mathcal{T} passed to the DISAGREEMENT or COUNTING tests as input corresponds to one member of \mathbb{C} which is constructed as follows: for each $t \in \mathcal{T}$ let (j_t, i_t, i'_t) be a corresponding paper, SB reviewer and DB reviewer assigned to this paper, then $\mathcal{C} = \bigcup_{t \in \mathcal{T}} (j_t, i_t, i'_t)$. Conversely, each member $\mathcal{C} \in \mathbb{C}$ gives rise to a family of sets of tuples $\mathbb{T}(\mathcal{C})$ which contains $2^{|\mathcal{C}|}$ elements and each element corresponds to a different allocation of reviewers in each triple $(j, i_1, i_2) \in \mathcal{C}$ to SB and DB conditions. For example, let $\mathcal{C} = \{(j, i_1, i_2), (j', i'_1, i'_2)\}$, then the family $\mathbb{T}(\mathcal{C})$ consists of four sets of tuples:

$$\begin{aligned} \mathcal{T}_1 &= \{(j, Y_{i_1 j}, X_{i_2 j}, w_j), (j', Y_{i_1 j'}, X_{i_2 j'}, w_{j'})\} \\ \mathcal{T}_2 &= \{(j, Y_{i_2 j}, X_{i_1 j}, w_j), (j', Y_{i_1 j'}, X_{i_2 j'}, w_{j'})\} \\ \mathcal{T}_3 &= \{(j, Y_{i_1 j}, X_{i_2 j}, w_j), (j', Y_{i_2 j'}, X_{i_1 j'}, w_{j'})\} \\ \mathcal{T}_4 &= \{(j, Y_{i_2 j}, X_{i_1 j}, w_j), (j', Y_{i_2 j'}, X_{i_1 j'}, w_{j'})\} \end{aligned}$$

Next, for concreteness assume that $\lambda \geq \mu$, that is, Algorithm 1 is used to construct a set \mathcal{T} . Then conditioned on the fact that the set of tuples \mathcal{T} constructed by the algorithm belongs to $\mathbb{T}(\mathcal{C})$, the randomness of the allocation of reviewers to conditions, the random assignment procedure used to assign reviewers to papers in each condition and randomness in the tie-breaking in the matching algorithm ensure that $\mathcal{T} \in \mathcal{U}[\mathbb{T}(\mathcal{C})]$, that is, all elements of $\mathbb{T}(\mathcal{C})$ are equally likely to be constructed and no other set of tuples can be constructed.

For each member $\mathcal{C} \in \mathbb{C}$, let $\mathbb{P}[\mathcal{C}]$ be probability that Algorithm 1 constructs a set of tuples that belongs to $\mathbb{T}(\mathcal{C})$. Notice that for some $\mathcal{C} \in \mathbb{C}$ we have $\mathbb{P}[\mathcal{C}] = 0$ which happens for example when $|\mathcal{C}| < n$, because Lemma 1 ensures that $|\mathcal{T}| = n$. Now, conditioning on any set \mathcal{C} with $\mathbb{P}[\mathcal{C}] > 0$ (instead of conditioning on A^*) and using Lemma 1 (instead of Observation B), we repeat the proof of Theorem 1 for both DISAGREEMENT and COUNTING tests. The unconditional result then follows from the law of total probability. The same argument applies to the case when $\lambda < \mu$ and hence we conclude the proof.

E.5.5 Proof of Corollary 2

The high-level idea of the proof is to construct matrices Π^{db} and Π^{sb} that simultaneously (for different choices of $\beta_0^{(\text{sb})}$ and $\beta_1^{(\text{sb})}$ coefficients) satisfy the null and the alternative hypotheses under the extended model given by equations (14a) and (14b).

We begin our construction from specifying values of $q_j, j \in [n]$. For each paper $j \in [n]$, let

$$q_j = \begin{cases} -1 & \text{if } w_j = 1 \\ 0 & \text{if } w_j = -1. \end{cases}$$

Then Π^{db} is generated from model (14a) with $\beta_0^{(\text{db})} = 0$ and $\beta_1^{(\text{db})} = 1$. In this way, for any reviewer $i \in [m]$ and for any paper $j \in [n]$, probability of acceptance $\pi_{ij}^{(\text{db})}$ satisfies:

$$M_0 : \log \frac{\pi_{ij}^{(\text{db})}}{1 - \pi_{ij}^{(\text{db})}} = q_j.$$

That is, for any reviewer $i \in [m]$ and for any paper $j \in [n]$ we have

$$\pi_{ij}^{(\text{db})} = \begin{cases} \frac{1}{1+e} & \text{if } w_j = 1 \\ 0.5 & \text{if } w_j = -1. \end{cases}$$

We now consider two different choices of coefficients for SB reviewers which result into two different models of behaviour of SB reviewers under the absence of bias:

$$M_1 (\beta_0^{(\text{sb})} = 1, \beta_1^{(\text{sb})} = 1) : \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = 1 + q_j$$

$$M_2 (\beta_0^{(\text{sb})} = 3/2, \beta_1^{(\text{sb})} = 2) : \log \frac{\pi_{ij}^{(\text{sb})}}{1 - \pi_{ij}^{(\text{sb})}} = \frac{3}{2} + 2q_j$$

Consider a matrix Π^{sb} whose components for each $i \in [m]$ and $j \in [n]$ are defined as follows:

$$\pi_{ij}^{(\text{sb})} = \begin{cases} 0.5 & \text{if } w_j = 1 \\ \frac{1}{1+e^{-1}} & \text{if } w_j = -1, \end{cases}$$

it is not hard to see that (i) entries of matrix Π^{sb} satisfy the model M_1 and (ii) for each paper $j \in \mathcal{J}$ corresponding entries of matrix Π^{sb} are larger than prescribed by model M_2 by $\delta > 0$ and for each paper $j \in \overline{\mathcal{J}}$ corresponding entries are smaller than those prescribed by M_2 by $\delta > 0$, where δ is some universal constant. Hence, depending on which model of SB reviewer under the absence of bias (M_1 or M_2) is correct, pair of matrices $(\Pi^{\text{db}}, \Pi^{\text{sb}})$ corresponds to the absence or presence of bias.

Given that matrices Π^{sb} and Π^{db} solely determine a distribution of reviewers' decisions, we have shown that reviewers' decisions are identically distributed under both null and alternative hypotheses under the extended version of the generalized logistic model. Hence, we conclude the proof by declaring that any algorithm that operates on reviewers' decision and keeps Type-I error below α must have power at most α under the alternative specified by models M_0, M_2 and matrices $\Pi^{\text{sb}}, \Pi^{\text{db}}$ for all values of $\min\{|\mathcal{J}|, |\overline{\mathcal{J}}|\}$ and hence violates the non-trivial power requirement.