Prior and Prejudice: The Novice Reviewers' Bias against Resubmissions in Conference Peer Review

IVAN STELMAKH, Carnegie Mellon University, USA

NIHAR B. SHAH, Carnegie Mellon University, USA

AARTI SINGH, Carnegie Mellon University, USA

HAL DAUMÉ III, University of Maryland, USA and Microsoft Research, USA

Modern machine learning and computer science conferences are experiencing a surge in the number of submissions that challenges the quality of peer review as the number of competent reviewers is growing at a much slower rate. To curb this trend and reduce the burden on reviewers, several conferences have started encouraging or even requiring authors to declare the previous submission history of their papers. Such initiatives have been met with skepticism among authors, who raise the concern about a potential bias in reviewers' recommendations induced by this information. In this work, we investigate whether reviewers exhibit a bias caused by the knowledge that the submission under review was previously rejected at a similar venue, focusing on a population of novice reviewers who constitute a large fraction of the reviewer pool in leading machine learning and computer science conferences. We design and conduct a randomized controlled trial closely replicating the relevant components of the peer-review pipeline with 133 reviewers (master's, junior PhD students, and recent graduates of top US universities) writing reviews for 19 papers. The analysis reveals that reviewers indeed become negatively biased when they receive a signal about paper being a resubmission, giving almost 1 point lower overall score on a 10-point Likert item ($\Delta = -0.78$, 95% CI = $[-1.30, -0.24]$) than reviewers who do not receive such a signal. Looking at specific criteria scores (originality, quality, clarity and significance), we observe that novice reviewers tend to underrate quality the most.

CCS Concepts: • **Human-centered computing** → **Empirical studies in collaborative and social computing**;

Additional Key Words and Phrases: peer review, biases, randomized controlled trial

ACM Reference Format:

Ivan Stelmakh, Nihar B. Shah, Aarti Singh, and Hal Daumé III. 2021. Prior and Prejudice: The Novice Reviewers' Bias against Resubmissions in Conference Peer Review. *Proc. ACM Hum.-Comput. Interact.* 5, CSCW1, Article 75 (April 2021), 17 pages. <https://doi.org/10.1145/3449149>

1 INTRODUCTION

In contrast to many other fields of science, where journals are the only established venues for research publication, in machine learning (ML) and computer science (CS), conferences are considered to be equally or even more attractive [72]. While being as selective as top journals, leading

Authors' addresses: Ivan Stelmakh, stiv@cs.cmu.edu, Carnegie Mellon University, School of Computer Science, 5000 Forbes Ave, Pittsburgh, Pennsylvania, USA, 15213; Nihar B. Shah, nihars@cs.cmu.edu, Carnegie Mellon University, School of Computer Science, 5000 Forbes Ave, Pittsburgh, Pennsylvania, USA, 15213; Aarti Singh, aarti@cs.cmu.edu, Carnegie Mellon University, School of Computer Science, 5000 Forbes Ave, Pittsburgh, Pennsylvania, USA, 15213; Hal Daumé III, hal3@umd.edu, University of Maryland, Department of Computer Science, College Park, Maryland, USA, 20742, Microsoft Research, 641 6th Ave, New York, New York, USA, 10011.

This work is licensed under a Creative Commons Attribution International 4.0 License.

© 2021 Copyright held by the owner/author(s).

2573-0142/2021/4-ART75. <https://doi.org/10.1145/3449149>

conferences ensure much shorter turnaround time thereby facilitating timely research dissemination and allowing authors to quickly resubmit their work to the next conference if it gets rejected. However, the explosion in the number of submissions received by top conferences has considerably challenged the sustainability of the conference peer-review process since the number of qualified reviewers is growing at a much slower rate [55, 56].

In an attempt to decrease the load on reviewers by discouraging resubmissions without substantial changes, several leading ML and CS conferences have started requesting or requiring authors to declare if a previous version of their submission was rejected at other peer-reviewed venues. For example, a top-tier conference in natural language processing EMNLP 2019 allowed authors to decide whether they want to disclose the past submission history and provide the summary of changes they made, making this information available only to senior committee members. A similar opportunity was offered at a leading conference in artificial intelligence and statistics (AISTATS 2017) with the exception that the information about past rejections was also available to regular reviewers. Additionally, the AISTATS 2017 conference implemented an automated review sharing with some past conferences, making these reviews visible to senior committee members after the initial reviewing was completed.

Other conferences make it mandatory for authors to disclose the past submission history: the NeurIPS conference — one of the most popular conferences in ML, which in 2019 received more than six thousand submissions — in 2020 required authors of previously rejected submissions to declare the changes they made to the current version of the paper. Another top conference in artificial intelligence (IJCAI 2020) went even further and made full reviews from past venues available to reviewers by requiring authors to include them in the submission file *before* the actual paper.

In addition to changes implemented by specific venues, the [openreview](#) platform — a growing conference management system that hosts the leading deep-learning conference ICLR and other forums — offers a novel approach towards managing conferences in a transparent manner by allowing organizers to make both accepted and rejected submissions accompanied with full reviews publicly available. This option has been used by the ICLR conference since 2017 and all reviews for papers submitted to the conference since then are now publicly available, allowing subsequent reviewers (even at different venues) of rejected papers to consult them at any time.

All these steps are supposed to facilitate improving the quality of peer review as reviews from past venues may reduce the burden on reviewers by allowing them to “quickly focus in on what previous issues were and how they may or may not have been fixed” [20, 36]. While being especially actual for the fields of ML and CS, similar discussion is also happening in other areas. For example, a survey conducted by Cals et al. [10] among editors of general medical journals showed that despite there being a concern that reviews from previous venues may bias subsequent reviewers, 45% of the participating editors prefer authors to indicate whether a paper has been previously rejected (the survey does not indicate how this information is supposed to be used in the review process). Only 24% of editors oppose this idea, and the rest are indecisive.

Despite the interest of journal editors and conference program committees in reusing reviews from past iterations of the review cycle, authors are less enthusiastic: ICFP — the ACM-sponsored conference on functional programming — in 2020 removed the option to upload the past reviews since no authors took this advantage in the previous edition of the conference. This skepticism suggests that authors do not believe that revealing old reviews will increase the acceptance chances of their submission.

Authors’ concerns are in fact supported by a long line of research in psychology that establishes susceptibility of human judgements to various biases [see 25, 29, 70, for overview], some of which are caused by additional (and sometimes irrelevant) information available to the decision-maker [6, 11,

18, 54, 70]. Projecting this evidence on the peer-review context, we hypothesize that the knowledge that a paper was rejected at a previous venue may negatively bias reviewers' evaluations, leading to what in this study we call a "resubmission bias". In other words, reviewers may judge the paper differently depending on whether they are notified about this paper being a resubmission or not.

To highlight the potential impact of the resubmission bias, we note that the peer-review process of the major machine learning conferences is far from being absolutely consistent and objective. Indeed, several controlled experiments [35, 47, 48] found very low degree of agreement between reviewers evaluating the same manuscript. This implies that the outcome of the review process heavily depends on the reviewers to whom the submission is assigned, introducing significant randomness in the final decisions. Hence, even a strong paper that deserves acceptance has a nontrivial chance of being rejected due to this randomness. The resubmission bias can *amplify this unfairness by putting a previously rejected paper at disadvantage in the subsequent conferences*. Given that success in academia is largely determined by the publication profile of a researcher, the resubmission bias can have far reaching consequences not just for a particular paper, but more generally also for career trajectories of researchers due to the widespread prevalence of the Matthew effect ("rich get richer") in academia [42, 58].

Therefore, in this work, we aim at testing for the presence of the resubmission bias in peer review. Focusing on the population of novice reviewers, in conjunction with a reviewer-recruiting process of ICML 2020 (a top ML conference), we design and conduct a randomized controlled trial to test the following research hypothesis:

Research Hypothesis: The signal about rejection from the same or similar venue in the past, received by novice reviewers before they read the paper, induces a bias in reviewers' evaluations.

In our hypothesis, we do not specify the direction of the effect and note that while authors are concerned about a negative bias, the bias can hypothetically be positive. For example, a reviewer may think that previously rejected papers have gone through another iteration of revisions and improvements and might be better on average than papers that have not previously been peer reviewed. Therefore, in this work we also aim to confirm the direction of the effect (if it exists).

Importantly, we caveat that in this study we target the population of novice reviewers and the findings we report in this paper must not be overgeneralized to the whole reviewer population. While the choice of the study participants is mostly justified by the difficulty of engaging senior reviewers in the experiment, we discuss below (Section 3) that novice reviewers constitute a large fraction of the leading ML and CS conferences reviewer pool.

In this work, we do not aim to support or undermine the idea of reusing the past reviews. Instead, the answer to our research question will inform conference organizers and journal editors about potential side-effects of reusing reviews, and will help them to carefully select the point during the peer-review pipeline (if at all) at which previous reviews become available to reviewers. For example, if the resubmission bias in peer review is present, then editors or program chairs may prefer to keep the past submission history hidden from reviewers in the initial stages of the process to ensure that the reviewers form an unbiased opinion about the paper first. Additionally, the results of our experiment can inform the current practices of novice reviewer training where more emphasis could be made on how to avoid the potential resubmission bias.

Past research on human decision making decisively establishes the presence of various cognitive biases in human judgement that are relevant to our research hypothesis. For example, the famous anchoring effect [70] manifests in human evaluations being dependent on an (irrelevant) piece of information received at the beginning of the decision task. We defer the discussion of the general literature on decision making and cognition to Section 5, but underscore that these findings do not necessarily transfer to our setting. Indeed, from the perspective of the dual-process model

of cognition [30, 59, 60], biases are properties of an autonomous heuristic system (System 1) whereas the demanding and rational reviewing task invokes the analytic system (System 2) that can potentially override such biases.

2 RELATED LITERATURE

Our paper contributes to a long line of empirical works studying various aspects of the academic peer-review process and in this section we discuss the most relevant papers.

Resnik and Smith [50], in the chapter of their book dedicated to peer review, discuss the presence of groupthink — a strong desire of group consensus that results in all deviating ideas being rejected — in peer review. One potential contribution to the overall groupthink effect is susceptibility of experts to social influence, as demonstrated by Teplitskiy et al. [66]. In their experiment, reviewers (faculty at US medical schools) first evaluated and rated submissions assigned to them, and then were exposed to scores presumably given to these submissions by other anonymous reviewers. Unbeknown to participants, these scores were randomly sampled to be either above or below their scores. The experiment demonstrated that 47% of reviewers decided to update their scores, and in all but one case (out of approximately 185) the update was in the direction of the external scores. This finding indicates a strong impact of social influence on experts' evaluations. The experiment of Teplitskiy et al. [66] was conducted alongside a review process of grant applications in which real awards were distributed; in this setup, reviewers are generally expected to reach a consensus, hence, updates of review scores do not necessarily indicate a bias in reviewers' judgements of the proposal quality and instead can be seen as attempts to decrease the inconsistency of evaluations. In our work, we design the experiment to remove the group deliberation component of peer review by eliminating the discussion stage of the process, and measure the bias in reviewers' attitude towards the submission.

Another widely documented bias [9, 15, 43, 53, and others] that manifests in peer review is the *positive-outcome bias* also known as the *file drawer problem*. This bias results in reviewers judging soundness of the experimental works depending on the outcome of the study: works in which the null hypothesis is rejected are rated more favourably than otherwise equal papers that show nonsignificant results.

One difference between the social influence, file drawer problem and resubmission bias which may result in different cognitive heuristics responsible for the effects is the stage of the review process at which the stimulus is received by reviewers. In case of the social influence, reviewers get biased by signals they receive *after* the review process and this bias can be attributed to the desire of consensus, attempt to improve the accuracy of a review or aim to achieve some other social goals [13, 50, 66]. The file drawer problem is induced by information observed *during* reviewing and may be attributed to the desire of accepting “newsworthy” (i.e., positive) studies for publications [9, 38]. Finally, the incarnation of the resubmission bias we study in this work completes the picture and is related to the information reviewers receive *prior* to reviewing submissions.

Confirmatory bias — a tendency of people to emphasize evidence that support their views and ignore or misinterpret those that do not — is another relevant effect connected to prior beliefs of reviewers. It was identified in context of peer review by Mahoney [39] who conducted an experiment in which reviewers were exposed to different versions of the manuscript with identical experimental procedure but different directions of the obtained results (either confirming or disproving the beliefs of reviewers). It turned out that reviewers who received the version contradicting their theoretical perspective were significantly harsher in their evaluations than those who received the version supporting their views. In Section 5 we draw further connections between the resubmission and confirmatory biases.

A separate line of work [8, 46, 68, and others] studies the presence of various biases, including gender, affiliation and fame biases in single-blind peer review (i.e., when author identities are visible to reviewers). The work of Tomkins et al. [68] is of particular interest as it identifies strong fame and affiliation biases¹ in reviewers' evaluations. As a result of this work, WSDM — a premier conference in web-inspired research — switched to double-blind reviewing (i.e., author identities are hidden from reviewers), demonstrating the potential impact of empirical research on peer-review practices.

Finally, our work contributes to the growing literature in computer science, both theoretical and empirical, that aims to understand and improve the conference peer-review process. These works develop methodologies to address various biases and other issues in peer review such as miscalibration [24, 52, 73], commensuration bias [45], strategic or dishonest behavior [3, 28, 63, 74], biases with respect to author demographics [40, 62, 68], and methods for better assignments of reviewers to papers [12, 17, 23, 32, 61]. We envisage that the biases discussed in the current paper may be mitigated by a combination of policy guidelines and such computational techniques.

3 EXPERIMENTAL SETUP

To test our research hypothesis of presence of the resubmission bias, we design a randomized controlled experiment that replicates the relevant components of the peer-review pipeline of machine learning conferences while giving us more control over the resubmission signal received by reviewers. Two main components of the peer-review process are pools of papers and reviewers, and we now describe how these pools were constructed.

Papers. We solicited $m = 19$ anonymized preprints in various sub-areas of machine learning. To ensure that participants of the experiment cannot obtain a signal about a paper being a resubmission from anywhere outside of the experimental context, we restricted the pool of papers to works that had not yet been accepted to any conference or journal and had never been submitted to conferences hosted by the openreview platform or any other venue that makes the list of rejected submissions publicly available. We note that in machine learning it is common to publish preprints on arXiv (arxiv.org) and we allow papers available on arXiv to be used in the experiment. Additionally, we allow papers that were previously presented at workshops — less formal and prestigious venues without proceedings — because it is also common to present a preliminary version of the work in a workshop and then submit the full version to the conference.

The final pool of papers consisted of working papers, papers under review, workshop publications and unpublished manuscripts. The papers were 6–12 pages long excluding references and appendices (a standard range for many ML and CS conferences) and were formatted in various popular journals' and conferences' templates with all explicit venue identifiers removed.

Reviewers. The reviewer pool of a typical machine learning conference consists of researchers at various seniority levels, working in areas covered by the conference. Perhaps unique to ML, a recent surge in the number of papers submitted to leading conferences has forced organizers to expand the reviewer pool by relaxing the qualification bar, that is, by introducing rather junior researchers to the reviewer pool: for example, 33% of the NeurIPS 2016 reviewers (1082 out of 3233) were graduate students [57]. Using data on the structure of the reviewer pool of the ICML 2020 conference, we estimate that approximately 35% of the reviewers are rather junior individuals who self-nominated and satisfied the screening requirements of having at least two papers published in some top ML venues and being a reviewer for at least one top ML conference in the past. Overall, we conclude that a significant fraction of ML reviewers are novices and in this study we concentrate on the population of novice and junior reviewers. We admit that in this work we do not approximate the

¹Meta-analysis they perform also reveals a bias against female-authored submissions.

general reviewer population by leaving out more experienced reviewers, but notice that at the ICML 2020 conference each submission was assigned to at least one reviewer from the aforementioned subset, making it a significant part of the reviewer community.

To recruit participants, we messaged mailing lists of five large large, top US universities (CMU, MIT, UMD, UC Berkeley and Stanford) and targeted master's and junior PhD students working in ML-related fields. The invitation also propagated to a small number of students outside of these schools through the word of mouth. The recruiting materials contained an invitation to participate in the ICML reviewer-selection experiment (selection of reviewers was indeed a key goal of the experiment and we studied the resubmission bias on top of it). Specifically, we asked participants to write a review for one paper and promised that those who provide a high-quality review will be invited to join the ICML 2020 reviewer pool. Being a reviewer at the flagship ML conference is a recognition of one's expertise and we expected that this potential benefit will motivate students to join our experiment. As a result, we received responses from $n = 200$ candidates, more than 90% of whom were master's and PhD students or recent graduates of the aforementioned universities, and all of them were added to the pool of participants without further screening. The research hypothesis we study in this work may be sensitive to awareness of participants; therefore, we employed deception and did not reveal the dual goal of this work to participants, that is, subjects were unaware that in parallel with selecting reviewers for ICML, we also want to measure the impact of the resubmission signal on the reviewers' attitude towards papers they review.

The experimental procedure closely followed the initial stages of the standard ML conference peer-review pipeline and was hosted using Microsoft Conference Management Toolkit (<https://cmt3.research.microsoft.com>). First, we asked participants to express their interest in reviewing specific papers by entering bids that take the following values: "Not Willing", "In a Pinch", "Willing" and "Eager". Thirteen participants did not enter any bids and were removed from the pool. The remaining participants were active in bidding (mean number of "Willing" and "Eager" bids is 4.7) and we assigned all of them to 1 paper each, where we tried to satisfy reviewer bids, subject to a constraint that each paper is assigned to at least 8 reviewers. As a result, 186 participants were assigned to a paper they bid either "Willing" or "Eager" and 1 participant was assigned to a paper they bid "In a Pinch" (this participant did not bid "Eager" or "Willing" on any paper).

Finally, we instructed participants that they should review the paper as if it was submitted to the real ICML conference with the exception that the relevance to ICML and the formatting issues (e.g., page limit, margins) should not be considered as criteria. To help participants in writing their reviews, we provided reviewer guidelines adapted from NeurIPS instructions (see supplementary materials) that discuss the best practices of reviewing. We gave participants 15 days to complete the review and then extended the deadline for 16 more days to accommodate late reviews as our original deadline interfered with the final exams at various US universities and the US holiday period.

Unbeknown to participants, we allocated half of them to the test condition and half to the control condition uniformly at random. Participants in the control condition did not receive the resubmission signal while subjects in the test condition were notified that the paper they are reviewing was rejected at the NeurIPS 2019 conference. Since the goal of the experiment declared to participants was to test out a new approach towards recruiting reviewers, the presence of the resubmission signal could have been confusing for the participants as this information is irrelevant to the task. To ensure that the presence of the signal does not look odd to reviewers, we incorporated it in a small author checklist placed on the first page of each submission as shown in Figure 1. Participants in the control condition received submissions with a single-item checklist asking about

Fig. 1. Different versions of checklists shown to reviewers in the test and control conditions. The answer to all questions were “Yes” for all papers in both conditions.

the code submission, while participants in the test condition were additionally informed that the paper is a resubmission.

4 ANALYSIS OF THE EXPERIMENT

Out of 187 participants who received a paper for review, 134 handed in the reviews (response rate of 71.7%). The remaining reviewers are excluded from subsequent analysis. Additionally, one of the participants misread the instructions and wrote a one-line review, rejecting the paper based on a minor violation of the standard ICML page limit. This participant was also excluded from the analysis as formatting violation was not considered as a review criterion in the experiment. Table 1 compares populations of reviewers between the test and control conditions using demographic information available to us.

The review form offered to participants (included in supplementary materials) was adapted from the NeurIPS 2019 form and contained 3 fields for open-ended feedback and 6 multiple choice questions in which reviewers were asked to give the overall score to the submission (10-point Likert item), evaluate the submission on 4 criteria (Originality, Quality, Clarity and Significance, 5-point Likert item for each criteria) and finally self-access their confidence in the scores assigned (5-point Likert item). Options offered for multiple choice questions were ordered from the most negative to the most positive and for the sake of analysis we associate these options to the numeric scales: 10-point scale for the overall score and 5-point scale for the other questions (larger numbers indicate more positive evaluations).² In the sequel of this section, we compare these values between the test and control reviewers to see whether the resubmission signal significantly impacted the behaviour of reviewers in the test condition.

	# PARTICIPANTS	#WITH PRIOR REVIEW EXPERIENCE	#WITH PUBLICATIONS
CONTROL	68	20	47
TEST	65	28	49

Table 1. Comparison of demographics of reviewers across the test and control conditions. All differences are not significant at the level $\alpha = 0.1$.

²We release numeric evaluations reported by reviewers and the dataset is available in supplementary materials.

4.1 Testing procedure and effect size

In this work, we employ a modification of the permutation test [19] that accounts for the fact that each paper is reviewed by different number of the test/control reviewers and is defined as follows. Recall that m stands for the number of papers in the experiment, for each paper $i \in [m]$ we let $\mathcal{R}_i^{\text{cntr}}$ (respectively $\mathcal{R}_i^{\text{test}}$) be a set of reviewers assigned to this paper in the control (respectively test) condition. Next, for any reviewer $j \in \mathcal{R}_i^{\text{cntr}}$ we use Y_i^j to denote a numeric evaluation given by reviewer j to paper i . Similarly, for any test reviewer $j \in \mathcal{R}_i^{\text{test}}$ we let X_i^j denote their evaluation of paper i . With this notation, we define the test statistic:

$$\tau = \frac{1}{m} \sum_{i \in [m]} \left[\underbrace{\frac{1}{|\mathcal{R}_i^{\text{test}}|} \sum_{j \in \mathcal{R}_i^{\text{test}}} X_i^j}_{M_i^{\text{test}}} - \underbrace{\frac{1}{|\mathcal{R}_i^{\text{cntr}}|} \sum_{j \in \mathcal{R}_i^{\text{cntr}}} Y_i^j}_{M_i^{\text{cntr}}} \right].$$

Our test statistic compares mean scores M_i^{test} and M_i^{cntr} received by each paper $i \in [m]$ from the test and control reviewers, respectively, and then averages differences of these scores across submissions. A negative value of the test statistic implies that reviewers in the test condition are more harsh than those in the control condition, whereas a positive value indicates that the test reviewers are more lenient.

Having defined the test statistic, we employ the permutation test to quantify the significance of its value. To this end, we permute reviewers within each paper between the test and control conditions uniformly at random, ensuring that the number of reviewers in each condition remains the same. We then recompute the value of the test statistic for 10,000 permutations and compare these values with the original value of the test statistic to obtain P -values.

In addition to reporting P -values of the test, we also provide the following measures of the effect size that capture different aspects of the effect:

- **Simple Effect Size** First, we note that our test statistic measures the between-conditions shift of the distributions of scores, that is, it represents the mean difference (in points of the corresponding Likert item) between scores given by two groups of reviewers. Hence, the value of the test statistic τ serves as a natural measure of the unscaled (simple) effect size [4]. Note that the larger the absolute value of the test statistic, the stronger the effect.
- **Scaled Effect Size (adaptation of Cohen's d)** Following recommendations of Cohen [14], we supplement the aforementioned simple effect size (which is expressed in the original units of analysis) with its scaled version, making it independent of the corresponding units. While the scaling inevitably obscures the interpretation of the resulting value, it helps to compare effect sizes between evaluations of the overall score (measured on the 10-point Likert item) and criteria scores (measured on 5-point Likert items). To report the scaled effect size, we use an adaptation of the Cohen's d [14]:

$$d = \frac{\tau}{\sqrt{v}},$$

where v is an upper bound on the variance of the terms M_i^{cntr} and M_i^{test} , $i \in [m]$, which is obtained by using the boundedness of numeric evaluations Y and X . Formally, $v = (k-1)^2/4$, where k is the number of points in the corresponding Likert item (10 for overall score and 5 for criteria scores). Note that we rely on the upper bound to avoid inflation of the effect size as terms corresponding to different papers have different variances, determined by the number of reviews written for this paper. The effect size computed in this manner is conservative because

it assumes the largest possible sample variance. Similar to the simple effect size, the larger the absolute value of the scaled effect size, the stronger the effect.

- **Relative Effect Size (stochastic superiority)** To reduce the impact of the extreme numeric evaluations on the effect size, we also consider the measure of stochastic superiority that is computed taking into account relative rather than absolute differences [71]:

$$A = \frac{1}{\sum_{i \in [m]} |\mathcal{R}_i^{\text{test}}| \times |\mathcal{R}_i^{\text{cntr}}|} \sum_{i \in [m]} \left(\sum_{j_1 \in \mathcal{R}_i^{\text{test}}} \sum_{j_2 \in \mathcal{R}_i^{\text{cntr}}} \left[\mathbb{I}(X_i^{j_1} > Y_i^{j_2}) + 0.5 \cdot \mathbb{I}(X_i^{j_1} = Y_i^{j_2}) \right] \right),$$

where $\mathbb{I}(\cdot)$ is an indicator function that equals 1 if its argument is correct and 0 otherwise. The denominator of this equation is the total number of (test, control) pairs of reviews written for the same paper. Each of these pairs contributes 1 to the numerator if the test review is more positive than the control review and 0.5 if the reviews contain the same numeric evaluation. Intuitively, this measure of the effect size equals to the empirical probability of a randomly sampled pair of (test, control) reviews written for the same paper having the test review more positive than the control (with ties broken uniformly at random). In contrast to the previous measures of the effect size, the further the value of A from 0.5, the stronger the effect.

In addition to reporting the point estimates of the effect sizes, we also report bootstrapped 95% confidence intervals computed over 5,000 iterations, where bootstrapping is performed at the level of papers. That is, in each iteration we create a new set of test (respectively, control) reviewers for each paper by sampling these reviewers with replacement from the set of the actual test (respectively, control) reviewers assigned to this paper. With all the preliminaries introduced, we are now ready to present the results of our experiment.

4.2 Results

Table 2 compares evaluations given by reviewers from the test and control conditions. The results indicate that reviewers in the test condition were significantly more strict in evaluating submissions, in average reporting almost 1 point lower overall score on the 10-point Likert item ($\Delta = -0.78$, 95% CI = $[-1.30, -0.24]$) than reviewers in the control condition. The difference appears to be small but considerable and we provide more discussion of the strength of the effect in the next section. Looking at the criteria score, we observe the similar trend of reviewers in the test conditions being stricter. Comparing unit-independent effect sizes (adaptation of Cohen's d and stochastic superiority), we note that the bias manifests the most in the evaluations of the submissions' quality which is known to be of high importance for the overall evaluation of the submission [45].

Overall, the data presented in Table 2 supports our research hypothesis, suggesting that junior reviewers are indeed susceptible to the resubmission bias. Notably, while the resubmission bias makes reviewers harsher, it does not seem to impact the confidence of reviewers. However, we qualify that self-evaluations of the confidence may not be a reliable measure of actual confidences, due to other biases manifesting in parallel. Indeed, as reviewers participate in the experiment to receive the invitation to join the pool of ICML reviewers, they may be reluctant to report too high or too low confidence as it may hypothetically hurt their chances. This hypothesis agrees with the observed data as 122 out of 133 reviewers reported confidence level 3 or 4 on the 5-point Likert item (recall that larger values indicate higher confidence).

As a final remark, we note that reported bootstrapped confidence intervals may be slightly more inaccurate than one would expect for the given sample size due specific nature of data. Indeed, 133 subjects of the experiment were broken into two groups (test and control) and distributed across 19

	P-VALUE	SIMPLE ES		SCALED ES		RELATIVE ES	
		SIZE	95% CI	SIZE	95% CI	SIZE	95% CI
OVERALL SCORE	.036*	-0.78	[-1.30, -0.24]	-0.17	[-0.29, -0.05]	0.42	[0.32, 0.52]
QUALITY	.005*	-0.46	[-0.69, -0.23]	-0.23	[-0.35, -0.11]	0.37	[0.27, 0.46]
CLARITY	.022*	-0.44	[-0.68, -0.19]	-0.22	[-0.34, -0.10]	0.43	[0.34, 0.52]
SIGNIFICANCE	.037*	-0.36	[-0.61, -0.10]	-0.18	[-0.30, -0.05]	0.43	[0.35, 0.50]
ORIGINALITY	.105	-0.21	[-0.40, -0.03]	-0.11	[-0.20, -0.02]	0.41	[0.32, 0.50]
CONFIDENCE	.902	-0.01	[-0.20, 0.17]	-0.01	[-0.10, 0.09]	0.50	[0.42, 0.59]

Table 2. Comparison of evaluations given by participants in the test and control conditions to papers assigned to them for review. All P -values are two-sided and are computed by the permutation test with 10,000 permutations. Asterisk indicates significance at the level $\alpha = 0.05$. Confidence intervals for effect sizes are computed using 5,000 bootstrapped samples.

papers. As mentioned above, the bootstrapping was performed at the level of papers. Thus, given that the sample size for each paper is small, the resulting intervals could capture the excessive variance or underestimate the actual variance. Nevertheless, the combination of different measures of the effect size and the results of the permutation test (which is guaranteed to control the false alarm probability even under the small sample size) suggest that the effect is present in the data.

5 DISCUSSION

The experiment we conduct in this work identifies the presence of the resubmission bias that manifests in junior reviewers being significantly harsher in evaluating submissions that they know were previously rejected from similar venues. The design of our experiment ensures the absence of unobserved confounders that could drive the result, as we obtain reviews for the *same* submission in both test and control conditions. In this section, we discuss some aspects of our experiment and suggest directions for future work.

Strength of the effect The size of the impact of the resubmission bias on the overall score received by submissions appears to be small according to the measures of the effect size we use in this work. We note, however, that top machine learning and computer science conferences are highly competitive and even small changes in reviewers' scores may have significant impact on the outcome of a submission and more generally on the researchers' career [58, 67]. As a concrete example, data from the ICML 2012 conference [34] demonstrates that papers with mean reviewer score 2.67 (on a 4-point Likert item) were 6 times more likely to be accepted than papers with mean score 2.33: the difference between these scores is a single point decrease in a single review. While this data is observational and does not account for potential unobserved "true quality" of submissions, it suggests that even small effects may result in large changes in the outcomes.

Next, we note that reviewers' evaluations are known to be subjective and have a high variance [5, 16, 31, 33, 39]. As a result, the effect that would be perceived as "strong" in a hypothetical within-subject experiment may be much less prominent in the between-subject design due to additional variance in evaluations due to subjectivity.

Finally, the pool of papers we have is diverse as it contains both papers that are likely to be accepted to the top conferences and papers that are not. In absence of the ground-truth ranking of papers and due to a limited sample size, we cannot look closer at the strength of the bias as a function of a paper quality, but past work suggests that some biases may be especially prominent in a subset of borderline papers [8]. As a direction for future work, it would be interesting to

(a) Distribution of the natural logarithm of the review length (in symbols). For clarity, we remove reviews shorter than 400 symbols from the set of reviews written by general ICML reviews and plot linearly-interpolated curves instead of discrete points.

(b) Mean rating of reviews given by area chairs across two groups of reviewers.

Fig. 2. Comparison of participants of the experiment with the general population of ICML reviewers.

understand the extent of the resubmission bias with a breakdown by the quality of the submission and some characteristics of reviewers (e.g., experience).

Population of participants The data obtained in this experiment unfortunately does not allow us to decisively align the population of the experiment participants with the general population of top-tier machine learning conference reviewers. Moreover, most of participants of the experiment would not be invited to join the reviewer pool of the ICML 2020 conferences through the standard ways of reviewer recruiting. However, as a result of the experiment, 52 participants whose reviews were found to be strong enough were invited to join the ICML reviewer pool. Therefore, these participants allow us to make some indirect comparisons to the general reviewer pool and we now present some relevant results. For a detailed analysis of performance of reviewers recruited through our experiment in the real ICML conference, we refer the reader to the companion paper [64].

Figure 2a juxtaposes the lengths of reviews written by different populations of reviewers at different venues. If we treat the length of a review as a measure of diligence, then we can see that reviewers selected in our experiment are not only more diligent than other participants of the experiment, but are also more diligent than the general reviewer population at the ICML conference. Perhaps surprisingly, even if we consider the participants of our experiment who did not get invited to review for the main conference, the length of their reviews stochastically dominates the length of the reviews written by general ICML reviewers with a significant margin, suggesting that they put a non-trivial effort in writing reviews.

Next, Figure 2b compares independent evaluations of review quality between the general population of ICML reviewers and reviewers recruited through our experiment. In the ICML review process each paper is assigned to several reviewers and one area chair. The area chairs are in charge of overseeing the reviewer activity, and at the end of the process, evaluate each reviewer on a 3-point Likert item: “Failed to Meet Expectations” (score 1), “Met Expectations” (score 2) and “Exceeded Expectations” (score 3). As seen in Figure 2b, experimental reviewers appear to produce

reviews of high quality according to the area chair's ratings, being on average better than general population of reviewers ($N_{\text{experimental}} = 111$, $N_{\text{general}} = 11624$, $\Delta = 0.18$, $P < .001$).³

As a word of caution, we note that comparisons we make in this section are based on observational data and may be influenced by confounding variables. For example, our experiment and the real conference employ different reviewer forms that could impact the length of the reviews. Additionally, experimental reviewers received at most 3 papers for review at ICML and only 1 paper in the experiment while general ICML reviewers received up to 6 submissions. There were also some other subtle differences between experimental and general reviewers in the assignment procedure of ICML that could contribute to the observed result.

Subject to the aforementioned caveats, indirect evidence presented in Figure 2 suggests that the participants of the experiment who got invited to join the ICML 2020 reviewer pool are comparable to the general population of top ML conference reviewers in terms of the length and quality of their reviews.

However, we underscore again that the findings of this work must be interpreted with care and must not be overgeneralized to a more senior population of reviewers. Indeed, a past study of reviewers' behaviour [66] demonstrated that seniority and expertise of reviewers impact their behaviour and hence serve as confounding factors. Therefore, we do not know whether the results reported in this paper extend to the whole reviewer pool, and an interesting direction for future work is to understand the presence of the resubmission bias in the general population of reviewers.

Impact of this study on the reviewer-selection process Recall that we conducted this study on top of the main experiment whose primary goal was to select participants to join the ICML reviewer pool. To minimize the effect of this study on the main experiment, prior to the selection process we removed numeric evaluations given by participants from the reviews, so the selection was based solely on the textual part of reviews. We then manually analyzed reviews that fall in the study team members' area of expertise, asked authors to comment on the review quality, and crowdsourced expert opinions for reviews that we were unable to evaluate ourselves. Combining feedback from these sources, we eventually invited 52 reviewers to join the ICML reviewer pool: 20 reviewers from the test condition and 32 reviewers from the control condition. The difference between the invitation rates appears to be insignificant at the level $\alpha = 0.05$ ($\Delta = 0.16$, $P = 0.075$).

Cognitive mechanism of the resubmission bias While in this work we do not aim to identify the cognitive mechanism of the resubmission bias, we now briefly comment on its relationship to several relevant cognitive biases known from general psychological literature.

- *Anchoring* [44, 65, 70] Human judgements and evaluations are known to depend heavily on an initial piece of information (anchor), even when the anchor is obviously irrelevant (e.g., a number generated by the wheel of fortune). The signal about the previous outcome received by reviewers may be seen as an anchor that prevents them from adjusting to a positive opinion about the paper; therefore, anchoring may be responsible for the resubmission bias.
- *Social proof* [2, 7, 13, 50] When being a part of a group, individuals are known to be susceptible to social influence, often exhibiting conformity to the opinion of the group. While in our experiment we removed the group deliberation component of the review process and participants were acting individually, they could consider the rejection from the previous venue as a decision made by a group of more experienced reviewers and decide to comply, deferring to the authority.
- *Confirmatory bias* [22, 37, 39, 49] Once having a belief about the state of the world, people tend to be selective in accepting new evidence: evidence that supports their views gets accepted more easily while contradicting evidence may be ignored or misinterpreted. In this paradigm,

³To evaluate significance, we use a permutation test treating the rating of each review as an independent random variable.

the resubmission bias can contribute to creating an initial belief that the paper under review is of a low quality that can later be exacerbated by a biased interpretation of strength and weaknesses of the submission.

- *Hindsight bias* [18, 27, 51] and *Outcome bias* [1, 6, 26, 41] Hindsight bias (“I knew it all along”) is a tendency of people to overestimate the predictability of the event after it becomes observed. Outcome bias is a slightly different effect that manifests in distortion of human judgements of decisions. When a stochastic outcome of the decision is observed, even having all the information available to the decision-maker at the time of the decision, people tend to judge the quality of the decision differently depending on whether they view the outcome as good or bad. While these effects do not directly correspond to the resubmission bias we study in this work, they suggest that availability of outcome information can adversely impact the evaluations.

Despite all of the above biases are known to be present in human judgements, their presence in peer review is not obvious because of the nature of the task performed by reviewers. Indeed, the reviewing task is analytical and requires rational thinking, thereby potentially reducing the reliance on heuristics responsible for cognitive biases [30, 59]. For example, Gino et al. [26], in context of the outcome bias, show that the outcome information biases participants less when they use rational mindset. Given that the task conducted by participants of the experiment of Gino et al. [26] is much less demanding than reviewing a paper, one could hypothesize that in peer review the cognitive heuristics may be overridden as reviewers naturally engage in rational and analytical thinking. However, our study identifies the presence of the resubmission bias in reviewers' judgements, demonstrating that the resubmission bias is strong enough to manifest even if reviewers put a non-trivial amount of cognitive effort into their work, as evidenced by the length of the reviews and expert judgements.

On general idea of reusing reviews In this work we do not aim to justify or oppose the idea of reusing reviews and do not say whether the resubmission bias is desirable or not. On the one hand, top-tier conferences are overloaded with substandard-quality papers making multiple rounds of submissions without major changes in a hope that they will get accepted at some point. Availability of past reviews could aid in identifying such submissions and optimizing the review efforts with respect to them, thereby reserving reviewers' effort for stronger papers. Additionally, when authors are aware of the resubmission bias, they may evaluate their work against higher standards and refrain from submitting works that are not yet in perfect shape to avoid the adverse effect on subsequent resubmissions, thereby further decreasing the load on reviewers.

On the other hand, the amount of randomness in the review process observed in several experiments [35, 47, 48] and in a multitude of anecdotal evidence (for example, see a survey of Nobel Prize laureates in economics by Gans and Shepherd 21) suggests that even strong papers that do not require revisions can be rejected by pure chance. It is also known [33, 69] that works that are novel or not mainstream, particularly those interdisciplinary in nature, face significantly higher difficulty in gaining acceptance. All these works can be put at significant disadvantage if reviewers are exposed to the resubmission signal as they may not have major aspects to be improved.

While in this work we do not resolve the aforementioned dichotomy, we identify the presence of the resubmission bias in reviewers' decisions and measure its effect size, letting the community and conference organizers decide on its desirability. Specifically, to perform the randomized controlled trial we choose perhaps the simplest format of exposing the previous outcome to reviewers, omitting the content of reviews received by submissions at past venues and withholding author statements that could explain changes (if any) made in the manuscript after the previous rejection.

In its simplicity, our stimulus can perhaps be seen as a point on a spectrum between two ways of exposure to the resubmission signal employed in the real world: the openreview system allows

subsequent reviewers to see that the paper was rejected and exposes any damning or hypercritical past reviews without letting authors present their view beyond what is stated in the discussion. In contrast, some other venues allow authors to present their case more clearly. Exploring the wider range of the aforementioned spectrum and identifying the desirable point on it (if any) is an interesting direction for future work.

Overall, the questions of whether the resubmission bias is desirable and how to reuse the reviews from the past venues in a fair and efficient way is an open question that requires a discussion in the community. In this work, we provide some concrete evidence that can help conference organizers to make informed decisions when designing the peer-review system. A promising direction for the future work is to evaluate the proposal of reusing reviews in a holistic manner, studying the interactions between different aforementioned positive and negative effects both theoretically and empirically.

ACKNOWLEDGMENTS

We thank authors of papers used in the experiment for contributing their works to this study. We are grateful to the support team of the Microsoft Conference Management Toolkit (CMT) who hosted our experiment and helped with many customization requests. We also thank authors of NeurIPS reviewer guidelines whose work we adapted to our setting. This study was approved by Carnegie Mellon University Institutional Review Board.

This work was supported in part by NSF CAREER Award CIF: 1942124 and NSF CIF: 1763734.

REFERENCES

- [1] Scott T. Allison, Diane M. Mackie, and David M. Messick. 1996. Outcome Biases in Social Perception: Implications for Dispositional Inference, Attitude Change, Stereotyping, and Social Behavior. In *Advances in Experimental Social Psychology*, Mark P. Zanna (Ed.). Vol. 28. Academic Press, 53 – 93. [https://doi.org/10.1016/S0065-2601\(08\)60236-1](https://doi.org/10.1016/S0065-2601(08)60236-1)
- [2] S. E. Asch. 1951. Effects of group pressure upon the modification and distortion of judgments. In *Groups, leadership and men; research in human relations*, H. Guetzkow (Ed.). Carnegie Press, 177–190.
- [3] Haris Aziz, Omer Lev, Nicholas Mattei, Jeffrey S Rosenschein, and Toby Walsh. 2019. Strategyproof peer selection using randomization, partitioning, and apportionment. *Artificial Intelligence* 275 (2019), 295–309.
- [4] Thom Baguley. 2008. Standardized or simple effect size: What should be reported? *British journal of psychology (London, England : 1953)* 100 (12 2008), 603–17. <https://doi.org/10.1348/000712608X377117>
- [5] Von Bakanic, Clark McPhail, and Rita J Simon. 1987. The manuscript review and decision-making process. *American Sociological Review* (1987), 631–642.
- [6] Jonathan Baron and John C. Hershey. 1988. Outcome bias in decision evaluation. *Journal of Personality and Social Psychology* 54, 4 (1988), 569–579.
- [7] R. S. Baron, J. A. Vandello, and B. Brunzman. 1996. The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology* 71, 5 (1996), 915–927.
- [8] Rebecca M Blank. 1991. The Effects of Double-Blind versus Single-Blind Reviewing: Experimental Evidence from The American Economic Review. *American Economic Review* 81, 5 (December 1991), 1041–1067. <https://ideas.repec.org/a/aea/aecrev/v81y1991i5p1041-67.html>
- [9] M. L. Callahan, R. L. Wears, E. J. Weber, C. Barton, and G. Young. 1998. Positive-outcome bias and other limitations in the outcome of research abstracts submitted to a scientific meeting. *JAMA* 280, 3 (1998), 254–257. <https://doi.org/10.1001/jama.280.3.254> arXiv:[1998.03.254](https://arxiv.org/abs/1998.03.254) [arXiv:/data/journals/jama/4568/jpv71042.pdf](https://arxiv.org/abs/1998.03.254)
- [10] J. W. Cals, C. D. Mallen, L. G. Glynn, and D. Kotz. 2013. Should authors submit previous peer-review reports when submitting research papers? Views of general medical journal editors. *Annals of family medicine* 11, 2 (2013), 179–181.
- [11] Thomas R Carretta and Richard L. Moreland. 1983. The Direct and Indirect Effects of Inadmissible Evidence1. *Journal of Applied Social Psychology* 13, 4 (1983), 291–309. <https://doi.org/10.1111/j.1559-1816.1983.tb01741.x>
- [12] L. Charlin and R. S. Zemel. 2013. The Toronto Paper Matching System: An automated paper-reviewer assignment system. <http://www.cs.toronto.edu/~zemel/documents/tpms.pdf>
- [13] Robert Cialdini and Noah Goldstein. 2004. Social Influence: Compliance and Conformity. *Annual review of psychology* 55 (02 2004), 591–621. <https://doi.org/10.1146/annurev.psych.55.090902.142015>
- [14] Jacob Cohen. 1992. A power primer. *Psychological Bulletin* 112, 1 (1992), 155–159.

- [15] G.B. Emerson, W.J. Warme, F.M. Wolf, J.D. Heckman, R.A. Brand, and S.S. Leopold. 2010. Testing for the presence of positive-outcome bias in peer review: A randomized controlled trial. *Archives of Internal Medicine* 170, 21 (2010), 1934–1939.
- [16] Edzard Ernst and Karl-Ludwig Resch. 1994. Reviewer bias: a blinded experimental study. *The Journal of laboratory and clinical medicine* 124, 2 (1994), 178–182.
- [17] T Fiez, N Shah, and L Ratliff. 2020. A SUPER* Algorithm to Optimize Paper Bidding in Peer Review. In *Conference on Uncertainty in Artificial Intelligence*.
- [18] Baruch Fischhoff and Ruth Beyth. 1975. I knew it would happen: Remembered probabilities of once–future things. *Organizational Behavior and Human Performance* 13, 1 (1975), 1 – 16. [https://doi.org/10.1016/0030-5073\(75\)90002-1](https://doi.org/10.1016/0030-5073(75)90002-1)
- [19] Ronald A. Fisher. 1935. *The design of experiments*. Oliver & Boyd, Oxford, England.
- [20] Paul Francis. 2008. Thoughts on Improving Review Quality. In *Proceedings of the Conference on Organizing Workshops, Conferences, and Symposia for Computer Systems* (San Francisco, California) (WOWCS'08). USENIX Association, USA, Article 7, 4 pages.
- [21] Joshua S. Gans and George B. Shepherd. 1994. How Are the Mighty Fallen: Rejected Classic Articles by Leading Economists. *Journal of Economic Perspectives* 8, 1 (March 1994), 165–179. <https://doi.org/10.1257/jep.8.1.165>
- [22] J. A. Garcia, Rosa Rodriguez-Sánchez, and J. Fdez-Valdivia. 2020. Confirmatory bias in peer review. *Scientometrics* 123 (2020), 517 – 533.
- [23] N. Garg, T. Kavitha, A. Kumar, K. Mehlhorn, and J. Mestre. 2010. Assigning Papers to Referees. *Algorithmica* 58, 1 (01 Sep 2010), 119–136. <https://doi.org/10.1007/s00453-009-9386-0>
- [24] Hong Ge, Max Welling, and Zoubin Ghahramani. 2013. A Bayesian model for calibrating conference review scores. <http://mlg.eng.cam.ac.uk/hong/nipsrevcal.pdf>
- [25] Thomas Gilovich, Daniel Kahneman, and Thomas Gilovich. 2002. Heuristics and Biases: The Psychology of Intuitive Judgment. (2002), 857.
- [26] Francesca Gino, Don Moore, and Max Bazerman. 2008. No Harm, No Foul: The Outcome Bias in Ethical Judgments. *Harvard Business School, Harvard Business School Working Papers* (01 2008). <https://doi.org/10.2139/ssrn.1099464>
- [27] Scott A. Hawkins and Reid Hastie. 1990. Hindsight: Biased judgment of past events after the outcomes are known. *Psychological Bulletin* 107, 3 (1990), 311–327.
- [28] Steven Jecmen, Hanrui Zhang, Ryan Liu, Nihar B. Shah, Vincent Conitzer, and Fei Fang. 2020. Mitigating Manipulation in Peer Review via Randomized Reviewer Assignments. In *NeurIPS*.
- [29] Daniel Kahneman. 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux, New York.
- [30] Daniel Kahneman and Shane Frederick. 2002. Representativeness revisited: Attribute substitution in intuitive judgment. *Heuristics and biases: The psychology of intuitive judgment* 49 (01 2002), 49–81. <https://doi.org/10.1017/CBO9780511808098.004>
- [31] Steven Kerr, James Tolliver, and Doretta Petree. 1977. Manuscript characteristics which influence acceptance for management and social science journals. *Academy of Management Journal* 20, 1 (1977), 132–141.
- [32] Ari Kobren, Barna Saha, and Andrew McCallum. 2019. Paper Matching with Local Fairness Constraints. In *ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*.
- [33] Michèle Lamont. 2009. *How professors think*. Harvard University Press.
- [34] John Langford. 2012. ICML acceptance statistics. <http://hunch.net/?p=2517> [Accessed: 05/30/2020].
- [35] Neil Lawrence and Corinna Cortes. 2014. The NIPS experiment. <http://inverseprobability.com/2014/12/16/the-nips-experiment>. [Accessed: 05/30/2020].
- [36] Hsuan-Tien Lin, Maria-Florina Balcan, Raia Hadsell, and Marc'Aurelio Ranzato. 2020. Getting Started with NeurIPS 2020. <https://medium.com/@NeurIPSConf/getting-started-with-neurips-2020-e350f9b39c28>. Accessed: 05/30/2020.
- [37] Charles G. Lord, Lee D. Ross, and Mark R. Lepper. 1979. Biased Assimilation and Attitude Polarization : The Effects of Prior Theories on Subsequently Considered Evidence. *Journal of Personality and Social Psychology* 37, 11 (1979), 2098–2109.
- [38] Joseph Lynch, Mary Cunningham, Winston Warme, Douglas Schaad, Fredric Wolf, and Seth Leopold. 2007. Commercially funded and United States-based research is more likely to be published; Good-quality studies with negative outcomes are not. *The Journal of bone and joint surgery. American volume* 89 (06 2007), 1010–8. <https://doi.org/10.2106/JBJS.F.01152>
- [39] Michael J Mahoney. 1977. Publication prejudices: An experimental study of confirmatory bias in the peer review system. *Cognitive therapy and research* 1, 2 (1977), 161–175.
- [40] Emaad Manzoor and Nihar B Shah. 2020. Uncovering Latent Biases in Text: Method and Application to Peer Review. In *INFORMS Workshop on Data Science*.
- [41] Greg W. Marshall and John C. Mowen. 1993. An Experimental Investigation of the Outcome Bias In Salesperson Performance Evaluations. *Journal of Personal Selling & Sales Management* 13, 3 (1993), 31–47. <https://doi.org/10.1080/08853134.1993.10753956>

- [42] Robert K Merton. 1968. The Matthew Effect in Science. *Science* 159 (1968), 56–63.
- [43] Ronald Moscati, Dietrich Jehle, David Ellis, Albert Fiorello, and Michael Landi. 1994. Positive-outcome Bias: Comparison of Emergency Medicine and General Medicine Literatures. *Academic Emergency Medicine* 1, 3 (1994), 267–271.
- [44] Thomas Mussweiler and Fritz Strack. 2001. Considering The Impossible: Explaining The Effects of Implausible Anchors. *Social Cognition - SOC COGNITION* 19 (04 2001), 145–160. <https://doi.org/10.1521/soco.19.2.145.20705>
- [45] Ritesh Noothigattu, Nihar B Shah, and Ariel D Procaccia. 2020. Loss Functions, Axioms, and Peer Review. In *ICML workshop on Incentives in Machine Learning*.
- [46] Kanu Okike, Kevin T. Hug, Mininder S. Kocher, and Seth S. Leopold. 2016. Single-blind vs Double-blind Peer Review in the Setting of Author Prestige. *JAMA* 316, 12 (09 2016), 1315–1316. <https://doi.org/10.1001/jama.2016.11014> arXiv:<https://jamanetwork.com/journals/jama/articlepdf/2556112/jld160026.pdf>
- [47] Elizabeth L. Pier, Markus Brauer, Amarette Filut, Anna Kaatz, Joshua Raclaw, Mitchell J. Nathan, Cecilia E. Ford, and Molly Carnes. 2018. Low agreement among reviewers evaluating the same NIH grant applications. *Proceedings of the National Academy of Sciences* 115, 12 (2018), 2952–2957. <https://doi.org/10.1073/pnas.1714379115> arXiv:<https://www.pnas.org/content/115/12/2952.full.pdf>
- [48] Eric Price. 2014. The NeurIPS experiment. <http://blog.mrtz.org/2014/12/15/the-nips-experiment.html>. [Accessed: 05/30/2020].
- [49] Matthew Rabin and Joel L. Schrag. 1999. First Impressions Matter: A Model of Confirmatory Bias. *The Quarterly Journal of Economics* 114, 1 (1999), 37–82.
- [50] David B. Resnik and Elise M. Smith. 2020. *Bias and Groupthink in Science's Peer-Review System*. Springer International Publishing, Cham, 99–113. https://doi.org/10.1007/978-3-030-36822-7_9
- [51] Neal J. Roese and Kathleen D. Vohs. 2012. Hindsight Bias. *Perspectives on Psychological Science* 7, 5 (2012), 411–426. <https://doi.org/10.1177/1745691612454303>
- [52] Magnus Roos, Jörg Rothe, and Björn Scheuermann. 2011. How to Calibrate the Scores of Biased Reviewers by Quadratic Programming. In *AAAI Conference on Artificial Intelligence*.
- [53] Robert Rosenthal. 1979. The file drawer problem and tolerance for null results. *Psychological Bulletin* 86, 3 (1979), 638–641.
- [54] Lee Ross, Mark Lepper, and Michael Hubbard. 1975. Perseverance in Self-Perception and Social Perception: Biased Attributional Processes in the Debriefing Paradigm. *Journal of personality and social psychology* 32 (12 1975), 880–92. <https://doi.org/10.1037/0022-3514.32.5.880>
- [55] D. Sculley, Jasper Snoek, and Alexander B. Wiltchko. 2019. Avoiding a Tragedy of the Commons in the Peer Review Process. *CoRR* abs/1901.06246 (2019). arXiv:[1901.06246](https://arxiv.org/abs/1901.06246) <http://arxiv.org/abs/1901.06246>
- [56] Nihar B. Shah. 2019. *Principled Methods to Improve Peer Review*. Retrieved May 27, 2020 from http://www.cs.cmu.edu/~nihars/publications/survey_peerreview_niharshah.pdf
- [57] Nihar B Shah, Behzad Tabibian, Krikamol Muandet, Isabelle Guyon, and Ulrike Von Luxburg. 2018. Design and analysis of the NIPS 2016 review process. *The Journal of Machine Learning Research* 19, 1 (2018), 1913–1946.
- [58] Flaminio Squazzoni and Claudio Gandelli. 2012. Saint Matthew strikes again: An agent-based model of peer review and the scientific community structure. *Journal of Informetrics* 6, 2 (2012), 265–275.
- [59] Keith Stanovich. 1999. *Who Is Rational? Studies of Individual Differences in Reasoning*. Mahwah, NJ: Erlbaum.
- [60] Keith Stanovich and Richard West. 2008. On the Relative Independence of Thinking Biases and Cognitive Ability. *Journal of personality and social psychology* 94 (05 2008), 672–95. <https://doi.org/10.1037/0022-3514.94.4.672>
- [61] Ivan Stelmakh, Nihar Shah, and Aarti Singh. 2018. PeerReview4All: Fair and Accurate Reviewer Assignment in Peer Review. *arXiv preprint arxiv:1806.06237* (2018).
- [62] Ivan Stelmakh, Nihar Shah, and Aarti Singh. 2019. On Testing for Biases in Peer Review. In *NeurIPS*.
- [63] I. Stelmakh, N. B. Shah, and A. Singh. 2020. Catch Me if I Can: Detecting Strategic Behaviour in Peer Assessment. In *ICML Workshop on Incentives in Machine Learning*.
- [64] Ivan Stelmakh, Nihar B. Shah, Aarti Singh, and Hal Daumé III. 2020. A Novice-Reviewer Experiment to Address Scarcity of Qualified Reviewers in Large Conferences. (2020). Preprint <http://www.cs.cmu.edu/afs/cs.cmu.edu/user/istelmak/www/icml/novice.pdf>.
- [65] Fritz Strack and Thomas Mussweiler. 1997. Explaining the Enigmatic Anchoring Effect: Mechanisms of Selective Accessibility. *Journal of Personality and Social Psychology* 73 (09 1997), 437–446. <https://doi.org/10.1037/0022-3514.73.3.437>
- [66] Misha Tepitskiy, Hardeep Ranub, Gary S. Grayb, Michael Meniettid, Eva C. Guinan, and Karim R. Lakhani. 2019. Social Influence among Experts: Field Experimental Evidence from Peer Review.
- [67] Stefan Thurner and Rudolf Hanel. 2011. Peer-review in a world with rational scientists: Toward selection of the average. *The European Physical Journal B* 84, 4 (2011), 707–711.
- [68] Andrew Tomkins, Min Zhang, and William D. Heavlin. 2017. Reviewer bias in single- versus double-blind peer review. *Proceedings of the National Academy of Sciences* 114, 48 (2017), 12708–12713. <https://doi.org/10.1073/pnas.1707323114>

- arXiv:<https://www.pnas.org/content/114/48/12708.full.pdf>
- [69] G David L Travis and Harry M Collins. 1991. New light on old boys: Cognitive and institutional particularism in the peer review system. *Science, Technology, & Human Values* 16, 3 (1991), 322–341.
- [70] Amos Tversky and Daniel Kahneman. 1974. Judgment under Uncertainty: Heuristics and Biases. *Science* 185, 4157 (1974), 1124–1131. <https://doi.org/10.1126/science.185.4157.1124>
arXiv:<https://science.sciencemag.org/content/185/4157/1124.full.pdf>
- [71] András Vargha and Harold D. Delaney. 2000. A Critique and Improvement of the CL Common Language Effect Size Statistics of McGraw and Wong. *Journal of Educational and Behavioral Statistics* 25, 2 (2000), 101–132. <https://doi.org/10.3102/10769986025002101>
- [72] George Vrettas and Mark Sanderson. 2015. Conferences versus journals in computer science. *Journal of the Association for Information Science and Technology* 66, 12 (2015), 2674–2684. <https://doi.org/10.1002/asi.23349>
- [73] Jingyan Wang and Nihar B Shah. 2019. Your 2 is My 1, Your 3 is My 9: Handling Arbitrary Miscalibrations in Ratings. In *AAMAS*.
- [74] Yichong Xu, Han Zhao, Xiaofei Shi, and Nihar Shah. 2019. On Strategyproof Conference Review. In *IJCAI*.

Received June 2020; revised October 2020; accepted December 2020