


## RESEARCH ARTICLE

## The role of author identities in peer review

Nihar B. Shah \*

Machine Learning Department and Computer Science Department, Carnegie Mellon University, Pittsburgh, Pennsylvania, United States of America

\* [nihars@cs.cmu.edu](mailto:nihars@cs.cmu.edu)

## Abstract

There is widespread debate on whether to anonymize author identities in peer review. The key argument for anonymization is to mitigate bias, whereas arguments against anonymization posit various uses of author identities in the review process. The Innovations in Theoretical Computer Science (ITCS) 2023 conference adopted a middle ground by initially anonymizing the author identities from reviewers, revealing them after the reviewer had submitted their initial reviews, and allowing the reviewer to change their review subsequently. We present an analysis of the reviews pertaining to the identification and use of author identities. Our key findings are: (I) A majority of reviewers self-report not knowing and being unable to guess the authors' identities for the papers they were reviewing. (II) After the initial submission of reviews, 7.1% of reviews changed their overall merit score and 3.8% changed their self-reported reviewer expertise. (III) There is a very weak and statistically insignificant correlation of the rank of authors' affiliations with the change in overall merit; there is a weak but statistically significant correlation with respect to change in reviewer expertise. We also conducted an anonymous survey to obtain opinions from reviewers and authors. The main findings from the 200 survey responses are: (i) A vast majority of participants favor anonymizing author identities in some form. (ii) The "middle-ground" initiative of ITCS 2023 was appreciated. (iii) Detecting conflicts of interest is a challenge that needs to be addressed if author identities are anonymized. Overall, these findings support anonymization of author identities in some form (e.g., as was done in ITCS 2023), as long as there is a robust and efficient way to check conflicts of interest.

 OPEN ACCESS

**Citation:** Shah NB (2023) The role of author identities in peer review. PLoS ONE 18(6): e0286206. <https://doi.org/10.1371/journal.pone.0286206>

**Editor:** Alberto Baccini, University of Siena, Italy, ITALY

**Received:** February 15, 2023

**Accepted:** May 8, 2023

**Published:** June 21, 2023

**Copyright:** © 2023 Nihar B. Shah. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

**Data Availability Statement:** The data pertaining to the survey is contained within the [Supporting information](#) files. The data pertaining to the peer-review process cannot be shared publicly because it comprises sensitive information about the reviewing. All the data pertaining to the reviewing in the ITCS 2023 conference is available with Dr. Yael Kalai ([yael@microsoft.com](mailto:yael@microsoft.com)) who was the program chair of the ITCS 2023 conference.

**Funding:** The author(s) received no specific funding for this work.

**Competing interests:** The authors have declared that no competing interests exist.

## Introduction

There is a lot of debate in various research communities on whether to anonymize author identities or not. The primary argument for anonymization is that it mitigates bias in the review process. Indeed, a number of experiments in computer science [1, 2] and in other fields [3–8] have found evidence that reviews are biased if author identities are visible to the reviewers. The most consistent finding across these studies is that reviews are biased favorably towards authors from higher-ranked affiliations or authors who are famous.

Traditionally, peer review in the field of theoretical computer science does *not* anonymize author identities. This topic has been hotly debated in the field with various opinions and discussions [9–13]. Some of the stated reasons against anonymization pertain to using author

identities to gain more confidence in hard-to-verify mathematical proofs or wacky ideas, more scrutiny for papers by authors with a history of bad science, ensuring similar papers are from different authors, mitigating incremental work by same authors, and allocate reviewers' limited time and effort accordingly.

This study is focused on the Innovations in Theoretical Computer Science (ITCS) 2023 conference. In computer science, conferences typically review full papers (and not just abstracts), are the terminal venue of publication, are rated at least at par with journals, and are usually competitive with 15–25% acceptance rates.

Given the aforementioned debate regarding anonymization of authors, in the review process of the ITCS 2023 conference, we took the following middle-ground approach. Author identities were initially hidden from reviewers. Once the reviewer submitted their review, they could see author identities. Reviewers were subsequently allowed to modify their review if they wished. The rationale behind this approach was that it could mitigate bias at least in the initial impression, and later allow reviewers to use author identities in a manner that has been posited in aforementioned discussions. Note that the authors were free to upload their (non anonymous) papers to preprint servers, social media, etc.

In this paper we address three key questions pertaining to author identities in the review process:

- *Identification*: Given that author identities were initially hidden, how many reviewers could identify the authors?
  - We address this question by asking reviewers to self-report their knowledge of author identities when they submit their review.
- *Usage*: Given that author identities were made visible to reviewers after they submitted initial reviews, how did the reviewers subsequently use the author identities to modify their reviews?
  - We address this question by analyzing the change in the review scores and reviewers' self-reported expertise after they are initially submitted.
- *Opinions*: What opinions or perspectives do participants have regarding anonymization of author identities?
  - We address this question via an anonymous survey of the participants in the ITCS 2023 review process.

We envisage these results to be useful in designing author-anonymization policies in an evidence-based fashion.

## Related work

A number of studies investigate whether revealing identities of authors to reviewers results in biases in the reviews. Perhaps the most well-known such experiment in computer science involves a randomized controlled trial in the Web Search and Data Mining (WSDM) 2017 conference [1] in data mining. This experiment found biases favoring authors who were famous or from top-ranked institutions. (See the paper [14] for some concerns regarding the experimental design and methods employed in [1].) The study [2] conducted an observational study that found biases favoring authors from top-ranked affiliations in the International Conference on Learning Representations (ICLR) conference in machine learning. A number of studies in other fields also find biases favoring authors who are famous or from top-ranked institutions [3–8]; see [15] for an overview on biases in peer review.

The studies [16, 17] survey researchers about their perceptions of the peer-review process. The surveys reveal that researchers across a number of fields prefer anonymizing author identities. In the adjoining field of information theory where anonymization is currently not the norm, a survey of the members of the community [18] shows significant support for trying out anonymization.

In addition to mitigating bias, another aspect of reviews is their “quality”. A randomized controlled trial between anonymized and non-anonymized reviewing found that there was no difference in the quality of reviews as assessed by editors and authors [19].

The problem of measuring possible bias in peer review has also spurred research on designing and evaluating the methods to do so. The aforementioned work [14] provides an experimental design framework, statistical tests, and associated theoretical guarantees for conducting such a measurement in a randomized controlled trial. The paper [2] uses difference-in-difference techniques to estimate bias exploiting the fact that the ICLR conference changed from not anonymizing to anonymizing author names in a certain year. In addition, the paper [2] also develops methods to test for biases based on the review text rather than just the review scores. The pair of papers [20, 21] debate the methods of computing the bias, with different metrics leading to different conclusions. Finally, the paper [22] focuses on the trade-off between running such a controlled experiment in the peer-review process and the efficiency of the peer-review process.

Even when author identities are not included in the submissions, there are ways in which they may be deciphered. First, a reviewer may have seen the paper outside of the review process such as on a preprint server, on social media, or in a talk. For instance, a little over half of the papers submitted to the NeurIPS 2019 conference were posted on arXiv, and among these papers, 21% were seen by at least one reviewer [23]. Second, some reviewers may actively search for their assigned papers online. For instance, the study [24] found that over a third of the reviewers in two top conferences searched for their assigned paper online. Third, the reviewer may be able to guess the identity of the authors based on the content of the paper. The study [25] asked reviewers in three conferences which anonymized author identities to guess the authors of papers. A total of 70%-86% of the reviews did not provide any guess (which could mean that the reviewer did not have any idea of the identity or that the reviewer chose to just not answer the question). Among the subset of reviews which contained a guess, 72%-85% guessed at least one author correctly. Some other works [26–28] investigate the possibility of identifying authors based on content by designing machine learning algorithms to predict the authors based on the paper’s content. In our experiment, we also investigate this aspect in ITCS 2023 by asking reviewers to self-report whether they can identify the authors. The paper [29] aims to identify the factors that are associated with successful anonymization in seven biomedical journals. They found no association with whether the journal had a long-standing policy of anonymized reviewing, but they did find that anonymization was less successful when the reviewer had a greater research experience.

Opening up author identities to reviewers near the end of the review process has been employed before in other computer science conferences [30]. There, the author identities are usually revealed right before the program committee meeting. Their stated reasons for this approach are that fully anonymous review might penalize a paper for failing to cite a certain prior work but this prior work may be unpublished work by the same authors, and anonymous review may not discover certain conflicts which the reviewer may know about. Our middle ground approach is also motivated by various other frequently mentioned uses of author identities, as discussed in the Introduction section. Most importantly, a key objective of this approach in ITCS 2023 was to understand the use of author identities by reviewers to enable subsequent policy-makers to design policies using this evidence.

An alternative policy that lies between non-anonymization and compulsory anonymization is that of not making anonymization compulsory, but allowing the authors the option to anonymize their papers. The paper [31] studies such a policy in a number of journals from IOP publishers, finding that authors of 22% of papers volunteered to anonymize, and furthermore, the highest-prestige authors anonymized less often but still substantially. This policy increased the acceptance rate of papers by “low-prestige” authors by 5.6%. Such a policy in the Nature group of journals had resulted in 12% of papers being anonymized for reviewing [32].

Questions about bias might arise more when decisions can be subjective, for instance, relying on criteria such as predictions of potential impact or novelty or interestingness to rank papers. There are various experiments in computer science and elsewhere that have found high disagreement among reviewers in terms of which paper is better [33–37], and more recently, experiments that have also found high disagreements among co-authors about their jointly authored papers [38]. Consequently, there have been various debates about the low acceptance rates and “selectivity” in computer science conferences [39–41].

Our work focuses on biases associated with author identities. In addition to this, there are a variety of other issues in peer review. See [42] for details. These issues lead to a variety of challenging theoretical problems which can lead to considerable practical impact if solved, but are not well explored in theoretical computer science.

## Analysis of reviews

In this section, we analyze the reviews to understand their dependence on author identities. We begin with a brief overview and statistics of the review process at ITCS 2023, and then present our analysis of the change in the reviews after the reviewer could see the author identities. (See <https://aspredicted.org/ng77i.pdf> for the preregistration of this analysis.)

### Overview and basic statistics of the review process

The ITCS 2023 conference received a total of 235 paper submissions. The set of reviewers comprised a program committee (PC) and external reviewers. The 40-member program committee were ultimately responsible for the acceptance/rejection decisions on all papers. The program committee also included the program chair of the conference. The members of the program committee could review papers themselves or ask experts not in the program committee, called external reviewers, to review. In ITCS 2023, there were 40 members in the program committee who did a total of 405 reviews. There were 315 external reviewers who did 361 reviews.

Participants were notified of the experiment before the review process started and were provided an option to decline participation in the experiment. (None of the participants declined). This work was reviewed and approved by the Carnegie Mellon University Institutional Review Board (IRB). The analysis pertained to reviews already provided in the conference and the IRB waived informed consent.

The conference asked authors to anonymize their identities on the submitted paper. The conference policy did not prohibit authors from posting their (non anonymous) papers online or on social media or give talks. Reviewers were not shown the authors’ identities until they submitted a review. After a reviewer submitted their review, they could see the authors’ identities for that paper, and could modify their review in any manner they chose. The reviewers were told of this policy beforehand, and were also told that the (changes in the) reviews will be analyzed.

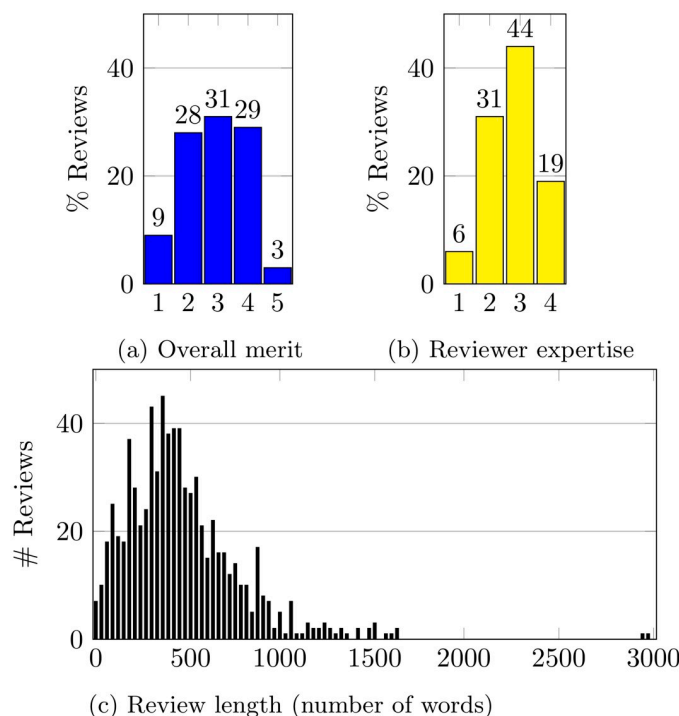
Reviewers were asked to fill out following form when submitting their review for any paper:

- Overall merit: 5. Strong accept, 4. Accept, 3. Weak accept, 2. Weak reject, 1. Reject

- Reviewer expertise: 4. Expert, 3. Knowledgeable, 2. Some familiarity, 1. No familiarity
- Paper summary \_\_\_\_\_
- Comments for authors \_\_\_\_\_
- Comments for PC \_\_\_\_\_
- Author identity: Please indicate your knowledge of author identities. [Visible to program chairs. Not visible to authors at any time.]
  - A. I know the author identities (e.g., saw a talk or on arxiv etc.)
  - B. I have a reasonable guess for author identities
  - C. I have no idea about author identities

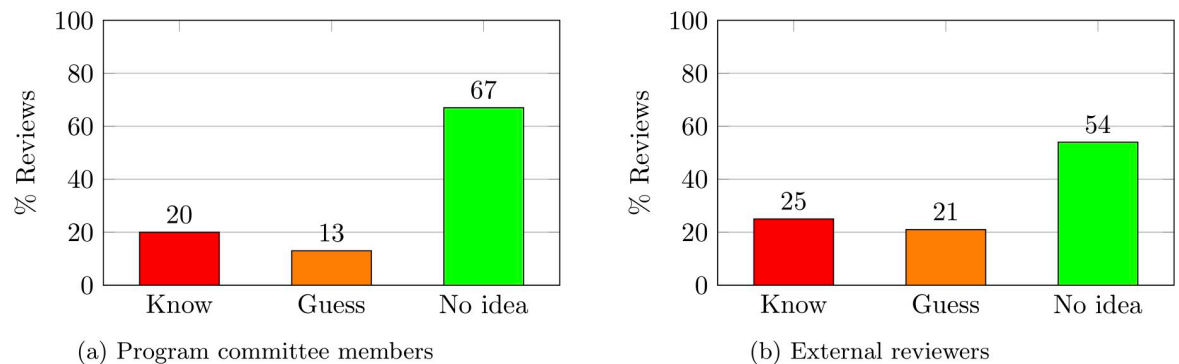
In Fig 1, we plot the basic statistics pertaining to the reviews—distribution of the overall merit scores, the self-reported reviewer expertise, and a histogram of the review length. Note that the reviewers could edit their reviews after they submitted an initial version, and these statistics pertain to the final state of each review. The mean length of the reviews was 480 words and median was 425 words.

We now come to the question about knowledge of author identities. It is important to note that this question was asked at the time the reviewer was submitting the review. This allows their answer to be based on a careful reading of the paper, and not just the title or abstract or a brief glance at the paper. In Fig 2 we plot the distribution of reviewers' answers to this question. We observe among both program committee members and external reviewers, **a majority of respondents report having no idea of the identities of the authors.**



**Fig 1. Basic statistics of the reviews in ITCS 2023.**

<https://doi.org/10.1371/journal.pone.0286206.g001>



**Fig 2. Distribution of reviewers' self reports about their knowledge of author identities.**

<https://doi.org/10.1371/journal.pone.0286206.g002>

### Change in reviews after initial submission

Recall that the reviewers could not see the identities of the authors of the paper initially, but the identities were made visible to them after they submitted a review for the paper. The reviewers could then edit their reviews, and in this section, we analyze the changes to the reviews after initial submission.

Before delving into the details, we must address one confounder. Once a program committee member submitted their review, they could also see all other reviews submitted for that paper. Thus any change in their review may be due to other reviews, and such an influence of other reviews is well documented in the literature [43]. Consequently, for program committee members, we restrict attention to only those reviews that were the first to be submitted for the paper. In other words, in this analysis, we only consider a review by a program committee member if there was no other review submitted for that paper when this review was first submitted. The external reviewers could not see other reviews at any time, and hence we consider all external reviews. This leaves us with 90 reviews by program committee members, along with the 361 reviews by external reviewers for our analysis. In the remainder of this section, we restrict attention to these 451 reviews.

We tabulate the percentage and number of reviews that changed in Table 1. In what follows, we provide a deeper analysis, where we also compute the correlation of the change in overall merit or reviewer expertise with the rank of the affiliations of the authors of the paper. We compute the rank of any affiliation based on [csrankings.org](https://csrankings.org) (2012 to 2022, Theory → Algorithms and Complexity). Affiliations that were not listed on [csrankings.org](https://csrankings.org) (such as industry labs) were ranked by the program chair. Then for any paper, we take the best rank among all affiliations associated with the authors of that paper. The p-values associated with the reported correlations are computed via the permutation test (with 10,000 permutations).

**Table 1. Number and percentage of reviews that changed after initial submission.**

|                          | Author identity | Change in overall merit | Change in reviewer expertise |
|--------------------------|-----------------|-------------------------|------------------------------|
| Program committee member | Know            | 1/20 = 5.0%             | 1/20 = 5.0%                  |
|                          | Guess           | 1/11 = 9.1%             | 0/11 = 0%                    |
|                          | No idea         | 4/59 = 6.8%             | 1/59 = 1.7%                  |
| External reviewer        | Know            | 5/90 = 5.6%             | 6/90 = 6.7%                  |
|                          | Guess           | 6/76 = 7.9%             | 4/76 = 5.2%                  |
|                          | No Idea         | 15/195 = 7.7%           | 5/195 = 2.6%                 |
| Total                    |                 | 32/451 = 7.1%           | 17/451 = 3.8%                |

<https://doi.org/10.1371/journal.pone.0286206.t001>

We first discuss the **change in overall merit scores**:

- We find that 7.1% reviews changed their overall merit scores whereas 92.9% left them unchanged (column 3 of [Table 1](#)).
- We compute Kendall's tau b correlation between the change in overall merit and the rank of the associated paper. A positive correlation means that the overall merit for a paper with a better rank increased after the initial submission. We find that the correlation coefficient is 0.019 and is not statistically significant ( $p = 0.7189$ ).
- Among the 32 reviews that changed overall merit, the mean change was  $-0.4$  and median change was  $-1$ . (Kendall's tau b correlation coefficient when restricted to these 32 changed reviews is 0.09.)

Next we discuss the **change in reviewer expertise**:

- We find that 3.8% reviews changed their overall merit scores whereas 96.2% left them unchanged (column 4 of [Table 1](#)).
- We investigate if the self-reported reviewer expertise increased if the paper was from a better-ranked affiliation and the reviewer initially gave a high overall merit score, or if the paper was from a lower-ranked affiliation and the reviewer initially gave a low overall merit score. In what follows, this would be implied by a positive correlation coefficient and the opposite by a negative correlation coefficient. We compute Kendall's tau b correlation between the affiliation-rank of the paper and the quantity (updated expertise—old expertise)\*(1 if overall merit = accept and -1 if overall merit = reject). We find that the correlation coefficient is 0.118 and is statistically significant at a level of 0.05 ( $p = 0.0265$ ).
- Among the 17 reviews that changed self-reported reviewer expertise, the mean change was  $+0.47$  and median change was  $+1$ . (Kendall's tau b correlation coefficient when restricted to these 17 changed reviews is 0.46.)

Also observe that among the aforementioned reviews that changed, a non-negligible fraction indicated that they knew the author identities, which suggests that either these reviews are changed for reasons not having to do with author identities, or that initially hiding the identities from the paper had some effect despite the reviewers' knowledge of author identities from elsewhere.

Finally, in 5 of the 315 reviews by external reviewers, the reviewer changed their answer about knowledge of author identities after initial submission. In two of these cases, they changed from 'no idea' to 'know', one changed from 'know' to 'no idea', one from 'guess' to 'no idea', and one from 'no idea' to 'guess'. None of the 90 reviews by program committee members changed the answer to the question on knowledge of author identities.

## Survey

We conducted an anonymous, optional survey among the participants of the ITCS 2023 conference. The survey was sent out via email on November 2, 2022 (soon after the paper acceptance decisions were announced) and was open till November 15, 2022. We first present the contents of the survey as shown to the participants, and then present aggregated responses.

## Questionnaire presented to participants

In what follows, we present the survey questionnaire verbatim as presented to the participants.

Our field has long debated whether to anonymize authors in the review process. The main reported benefit of anonymizing authors is that of reducing (biased) dependence on author identities in the review, as has been found in a number of experiments in other fields. A number of cons are also often discussed:

- Challenges in conflict-of-interest detection
- Reviewers want to use author identities for some part of their review (e.g., using track record to gain confidence in whacky ideas or complex proofs)
- Reduces ambiguity w.r.t. prior literature (e.g., can know that a preprint was by the same authors)
- Imperfectness of anonymization.

ITCS 2023 adopted a middle ground where reviewers initially did not see author identities, but could see after submitting their initial reviews, and could then modify their reviews. We would love to know your thoughts on this debate of anonymizing authors. When answering, please assume that conflict-of-interest detection is taken care of even if authors are anonymized (as done in many other conferences that adopt author anonymization).

I participated in ITCS 2023 as:

- PC member
- External reviewer
- Author

Preferences regarding anonymizing authors: (The ordering of the first four options was randomized to be either this order or reverse order).

- I prefer anonymizing author identities through the entire review process
- I prefer anonymizing author identities until the time for PC meetings/discussions
- I prefer anonymizing author identities until the reviewer submits their initial review
- I prefer not anonymizing author identities
- Other (please specify) \_\_\_\_\_

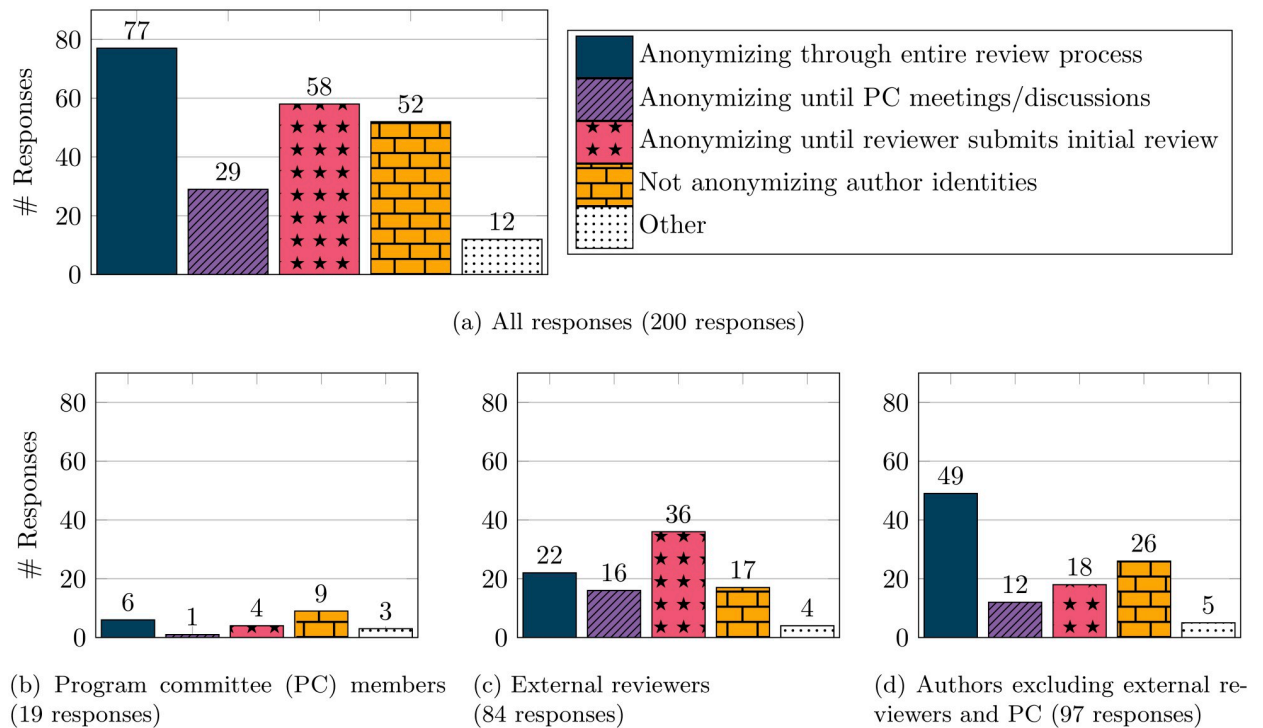
Please share the reasons for your preferences, as well as any other comments/experiences/opinions on this topic: \_\_\_\_\_

### Analysis of survey responses

We received 200 responses to the questionnaire. In this section, we summarize and analyze these responses. The full set of responses is available as [S1 Data](#).

**Quantitative analysis.** We begin with an analysis of the responses to the quantitative question on the respondents' preferences about anonymizing authors. We report the results in [Fig 3](#). (Note that participants were allowed to choose more than one option. Here we report the total counts in terms of the number of respondents who chose any option, hence the sum of all options in [Fig 3](#) is greater than 200. An alternative approach would be as follows: If a participant chooses  $k$  options, then count each chosen option as  $\frac{1}{k}$ . The result from this alternative approach is qualitatively similar to the current results, and hence we omit it for brevity.) The key takeaway is that there is considerable support for anonymizing authors in at least some part of the review process, particularly from participants outside the program committee.





**Fig 3. Answers given by respondents to the prompt “Preferences regarding anonymizing authors”.**

<https://doi.org/10.1371/journal.pone.0286206.g003>

Notice in Fig 3(a) that twelve respondents chose the option ‘other.’ These respondents also provided comments alongside this choice. The three PC members who selected ‘other’ also indicated support for anonymization in some form, such as anonymization to reviewers but not to program committee members. Five of the twelve respondents who selected ‘other’ and did not select any other option. Three of these five were supportive of anonymizing authors in peer review, but were concerned about challenges in detecting conflicts of interest (CoI). In absence of other ways for CoI detection, they supported anonymizing except for CoI detection. One of them strongly supported anonymization at least for external reviewers. One other respondent did not have an opinion. The remaining seven respondents who selected ‘other’ also selected other options alongside, and are already counted in Fig 3(a). We discuss their text comments below along with all other general comments made by respondents.

**Free-text comments.** We finally discuss the free-text comments provided by respondents. A total of 103 respondents provided free-text comments, and in what follows we summarize these comments. We saw in Fig 3 that there were disagreements among respondents on the best policies for anonymizing author identities, and these disagreements are also reflected in the free-text comments.

We begin with **comments that were common to a large number of respondents.**

- Respondents opine that revealing author identities can bias the reviews.
- Respondents appreciated the initiative taken by ITCS 2023 in adopting a middle-ground approach.
- Respondents complained about problems in avoiding conflicts-of-interest when assigning external reviewers.

We now discuss the remaining comments, each of which was given by **one or few respondents**. We have broadly classified these comments in terms of their implications regarding author anonymization.

• **Problems with anonymization / benefits of knowing author identities:**

- Reviewers can get biasing information about authors, e.g., from preprints uploaded by authors or from the English writing style in the submission.
- Anonymization would create weird incentives around posting papers on preprint servers, where authors might delay posting papers to stay anonymous or post earlier to make their identities public, depending on what they think would benefit them.
- May subsequently lead to undesirable policies such as preprint and talk embargoes.
- It can be challenging for authors to ensure that they do not inadvertently identify themselves.
- Anonymization makes rejecting a paper of a rival easier.
- Anonymization signals mistrust in the integrity of the community.
- Authors may strongly criticize their own past work to bolster the current submission, or may make only incremental contributions without being noticed under anonymization.
- Student authored papers should have lower bars, or can be shepherded instead of outright rejection.
- Some authors have a history of making errors or submitting bad papers, so knowing author identities can help assess correctness concerns.
- It is hard to assess quality in real time, and author reputation can serve as a good distinguisher, for instance, if the reviewer is not entirely familiar with the paper or if there are very long proofs.
- Reviewers should know who the authors are eventually, to give due credit if they learn something during the review process that they later use.
- If there are multiple closely related papers on a subject, it is important to understand to what extent these are distinct groups of authors.
- Need to specify anonymization policies for various situations, e.g., submitting a code repository.

• **Benefits of anonymizing:**

- Anonymization welcomes new authors and helps bring new ideas.
- Ideas and proofs should talk for themselves, and authors should explain new ideas well. Author name should not be a substitute.
- It will force reviewers to read the paper carefully rather than judging the quality based on authors.
- It is for the papers in the grey area that author identity will bias decisions, and these papers are precisely the ones for which the discussion is critical.
- Revealing author identities may make the reviewer more careful if the reviewer knows that the identities will be known at the discussion phase.

- **Benefits of middle ground:**
  - Prevents initial bias but then helps reap benefits of not anonymizing, thereby achieving both goals.
  - People get curious and will look up the paper, but they will be patient enough to wait until after the first review if they know they will get a chance to then see and potentially update their review.
  - Does not do any harm (presuming the chair can see the changes made).
- **Problems with middle ground:**
  - Puts back the bias into the process.
- **Suggestions:**
  - The program committee members should be able to see the name in the online peer-review platform but not on the paper.
  - Ask the reviewers to give two scores, one on the quality of the paper/result/novelty and another on confidence in the correctness/technical part. Among them, only the “correctness confidence score” can be modified after seeing the authors’ names after submitting an initial review.
  - It should be the reviewer’s choice as to whether they want author identities anonymized or not.
  - Prohibiting arXiv and other external dissemination is not good, but there is no harm in anonymizing within the review process.
  - Can have a two-tiered program committee where the senior members are aware of the author identities but do not actively participate in reviewing, but helps to resolve any issues with conflicts of interest.
  - In the middle ground, program chairs should be able to see any change in the reviews once the reviewer sees the author identities.

## Discussion and limitations

Our results suggest support for some form of anonymizing authors in peer review. Doing so can help mitigate biases pertaining to initial impressions about authors. In the literature, there are some stated benefits of not anonymizing authors, such as using author identities to gain more confidence in hard-to-verify mathematical proofs or wacky ideas, more scrutiny for papers by authors with a history of bad science, ensuring similar papers are from different authors, mitigating incremental work by same authors, etc. While these stated benefits are debated in the survey responses, if a venue wishes to realize them, the approach followed by ITCs 2023 may allow them to do so. Although this may result in bias appearing in the process after revelation of author identities, in ITCs 2023, we did not find much change in the overall merit scores after author identities were revealed. (Note that we reported being unable to check changes in review text or eventual discussions, leaving open the possibility of author identities playing a role there.) It is also sometimes stated that anonymizing authors is not useful as author identities may be guessed from the contents of the paper. However, we found that a majority of reviewers were unable to guess the authors’ identities.

Another stated benefit of knowing author identities is the ease of checking conflicts of interest. This was indeed a challenge at ITCS 2023, where several participants in the survey complained about problems in ensuring that papers are assigned to external reviewers who do not have conflicts of interest with the paper. Thus an anonymization of authors should be accompanied by an efficient and rigorous process to check conflicts of interest. One way to do so is to have an automated system to check for co-authorship and affiliation conflicts, as is done in machine learning conferences. An alternative “manual” option is to have a small set of volunteers who can check conflicts for any external reviewer that a program committee member wishes to invite for reviewing. A third possible solution, which leads to some reduction in anonymity, is to make author identities anonymous only to external reviewers but not program committee members.

Note that importantly, the conference did not impose any restrictions on authors regarding posting their (non anonymous) papers elsewhere such as on preprint servers. Such policy choices are widely debated [24] in some other research communities, where arguments in favor of such restrictions point to ensuring more thorough anonymization, whereas arguments against such restrictions include allowing free and open dissemination of research and furthermore of unintentionally biasing against the very people that it is supposed to protect [42, Chapter 7]. The policies of ITCS 2023, as well as the analysis and discussion in the present paper pertain to the absence of such restrictions.

This study also has several limitations that we discuss below.

- As briefly mentioned above, we did not have access to any changes in the text of the reviews. (Our preregistration included analysis of the change of text. However, we could not access the logs of the review text from the HotCRP conference management platform on which the peer-review process was conducted.) Author identities may also have played a role in the discussions between program committee members but we do not have logs of these discussions to draw any inference.
- One manner in which reviewers may use author identities is to gain confidence in their evaluation. We attempt to measure this, however, the review questionnaire asked reviewers to self-report their expertise. The confidence and expertise of any reviewer may be related, but not identical. For instance, a self-report of expertise may pertain to the topic of the paper whereas a self-report of confidence may pertain to their evaluation of the paper.
- The survey we conducted was anonymous, and hence comes with the usual caveats associated with anonymous surveys such as possible selection biases. We do stratify the responses by the respondents self-reported role in ITCS 2023 (program committee member, external reviewer, or author) but investigation of selection biases or stratification along other attributes is not possible.
- The experiment was announced to reviewers prior to them starting to review, and they were told that the logs of the review revisions will be analyzed. It is thus possible that this information may have changed the reviewer’s behavior (Hawthorne effect): the reviewer may be hesitant to change their reviews after seeing the author identities thinking that a change in their review could be undesirably perceived as tied to knowledge of author identities.

As for the last point above, we think that future conferences adopting such policies of partial anonymization should consider imparting transparency with respect to dependence on author identities, in which the program chairs, or the program committee, or even the authors can see the revisions of any review.

All in all, we hope that this investigation will lead to more discussion and evidence-based policy design for an improved peer-review process.

## Supporting information

**S1 Data. Responses to the survey.**  
(CSV)

## Acknowledgments

We sincerely thank the ITCS 2023 program chair Yael Kalai for trying out this middle ground approach to author anonymization, for facilitating the analysis, and for valuable inputs throughout the analysis. We also thank Pravesh Kothari for very helpful discussions. We are grateful to all the program committee members and reviewers of ITCS 2023 for their efforts in the review process, and all the respondents of our survey for sharing their opinions and suggestions. Finally, we thank the two anonymous reviewers for their valuable feedback on the manuscript and the Academic Editor Alberto Baccini for efficient handling of the manuscript.

## Author Contributions

**Conceptualization:** Nihar B. Shah.

**Data curation:** Nihar B. Shah.

**Formal analysis:** Nihar B. Shah.

**Investigation:** Nihar B. Shah.

**Methodology:** Nihar B. Shah.

**Project administration:** Nihar B. Shah.

**Resources:** Nihar B. Shah.

**Software:** Nihar B. Shah.

**Validation:** Nihar B. Shah.

**Visualization:** Nihar B. Shah.

**Writing – original draft:** Nihar B. Shah.

**Writing – review & editing:** Nihar B. Shah.

## References

1. Tomkins A, Zhang M, Heavlin WD. Reviewer bias in single-versus double-blind peer review. *Proceedings of the National Academy of Sciences*. 2017; 114(48):12708–12713. <https://doi.org/10.1073/pnas.1707323114> PMID: 29138317
2. Manzoor E, Shah NB. Uncovering Latent Biases in Text: Method and Application to Peer Review. In: AAAI conference on Artificial Intelligence; 2021.
3. Okike K, Hug KT, Kocher MS, Leopold SS. Single-blind vs double-blind peer review in the setting of author prestige. *JAMA*. 2016; 316(12):1315–1316. <https://doi.org/10.1001/jama.2016.11014> PMID: 27673310
4. Blank RM. The effects of double-blind versus single-blind reviewing: Experimental evidence from the American Economic Review. *The American Economic Review*. 1991; p. 1041–1067.
5. Ross JS, Gross CP, Desai MM, Hong Y, Grant AO, Daniels SR, et al. Effect of blinded peer review on abstract acceptance. *JAMA*. 2006; 295(14):1675–1680. <https://doi.org/10.1001/jama.295.14.1675> PMID: 16609089
6. Garfunkel JM, Ulshen MH, Hamrick HJ, Lawson EE. Effect of institutional prestige on reviewers' recommendations and editorial decisions. *JAMA*. 1994; 272(2):137–138. <https://doi.org/10.1001/jama.1994.03520020063017> PMID: 8015125

7. Fisher M, Friedman SB, Strauss B. The effects of blinding on acceptance of research papers by peer review. *JAMA*. 1994; 272(2):143–146. <https://doi.org/10.1001/jama.1994.03520020069019> PMID: 8015127
8. Huber J, Inoua S, Kerschbamer R, König-Kersting C, Palan S, Smith VL. Nobel and novice: Author prominence affects peer review. *Proceedings of the National Academy of Sciences*. 2022; 119(41): e2205779119. <https://doi.org/10.1073/pnas.2205779119> PMID: 36194633
9. Fortnow L. Flying Blind; 2018. <http://blog.computationalcomplexity.org/2018/02/flying-blind.html>.
10. Barak and other commentators B. On double blind reviews in theory conferences; 2018. Windows On Theory blog <https://windowsontheory.org/2018/01/11/on-double-blind-reviews-in-theory-conferences/>.
11. Venkatasubramanian S. Double blind review at theory conferences: More thoughts; 2018. <http://blog.geomblog.org/2018/01/double-blind-review-at-theory.html>.
12. Reingold O. Just defined bias? Not Sure; 2018. <https://theorydish.blog/2018/01/25/just-defined-bias-not-sure/>.
13. Mitzenmacher M. Double-blind, ALENEX; 2018. <http://mybiasedcoin.blogspot.com/2018/01/double-blind-alenex.html>.
14. Stelmakh I, Shah N, Singh A. On Testing for Biases in Peer Review. In: *NeurIPS*; 2019.
15. Haffar S, Bazerbachi F, Murad MH. Peer review bias: a critical review. In: *Mayo Clinic Proceedings*. vol. 94. Elsevier; 2019. p. 670-676.
16. Ware M. Peer review: benefits, perceptions and alternatives. *Publishing Research Consortium*. 2008;4:1-20.
17. Mulligan A, Hall L, Raphael E. Peer review in a changing world: An international study measuring the attitudes of researchers. *Journal of the Association for Information Science and Technology*. 2013; 64(1):132–161.
18. Shah N, Draper S. IEEE Information Theory Society Report on the 2021 Diversity and Inclusion Survey; 2022. [https://www.itsoc.org/sites/default/files/2022-04/ITSoc\\_D\\_I\\_Survey\\_Analysis.pdf](https://www.itsoc.org/sites/default/files/2022-04/ITSoc_D_I_Survey_Analysis.pdf).
19. Justice AC, Cho MK, Winker MA, Berlin JA, Rennie D, Investigators P, et al. Does masking author identity improve peer review quality?: A randomized controlled trial. *Jama*. 1998; 280(3):240–242. <https://doi.org/10.1001/jama.280.3.240> PMID: 9676668
20. Madden S, DeWitt D. Impact of double-blind reviewing on SIGMOD publication rates. *ACM SIGMOD Record*. 2006; 35(2):29–32. <https://doi.org/10.1145/1147376.1147381>
21. Snodgrass R. Single-versus double-blind reviewing: an analysis of the literature. *ACM Sigmod Record*. 2006; 35(3):8–21. <https://doi.org/10.1145/1168092.1168094>
22. Jecmen S, Zhang H, Liu R, Fang F, Conitzer V, Shah NB. Near-Optimal Reviewer Splitting in Two-Phase Paper Reviewing and Conference Experiment Design. In: *Conference on Human Computation and Crowdsourcing (HCOMP)*; 2022.
23. Beygelzimer A, Fox E, d'Alché Buc F, Larochelle H. What we learned from NeurIPS 2019 data; 2019. Available from: <https://neuripsconf.medium.com/what-we-learned-from-neurips-2019-data-111ab996462c>.
24. Rastogi C, Stelmakh I, Shen X, Meila M, Echenique F, Chawla S, et al. To ArXiv or not to ArXiv: A Study Quantifying Pros and Cons of Posting Preprints Online. *arXiv preprint arXiv:220317259*. 2022;.
25. Le Goues C, Brun Y, Apel S, Berger E, Khurshid S, Smaragdakis Y. Effectiveness of anonymization in double-blind review. *Communications of the ACM*. 2018; 61(6):30–33. <https://doi.org/10.1145/3208157>
26. Hill S, Provost FJ. The myth of the double-blind review? Author identification using only citations. *SIGKDD Explorations*. 2003; 5:179–184. <https://doi.org/10.1145/980972.981001>
27. Caragea C, Uban A, Dinu LP. The myth of double-blind review revisited: ACL vs. EMNLP. In: *EMNLP-IJCNLP*; 2019. p. 2317-2327.
28. Matsubara Y, Singh S. Citations Beyond Self Citations: Identifying Authors, Affiliations, and Nationalities in Scientific Papers. In: *International Workshop on Mining Scientific Publications*; 2020. p. 9-20.
29. Cho MK, Justice AC, Winker MA, Berlin JA, Waeckerle JF, Callahan ML, et al. Masking author identity in peer review: What factors influence masking success? *JAMA*. 1998; 280(3):243–245. <https://doi.org/10.1001/jama.280.3.243> PMID: 9676669
30. Hicks M. POPL 2012: FAQ concerning double-blind reviewing.; 2012. <http://www.cs.umd.edu/~mwh/dbr-faq.html>.
31. Smirnova I, Romero DM, Teplitskiy M. The bias-reducing effect of voluntary anonymization of authors' identities: Evidence from peer review; 2023. <https://papers.ssrn.com/abstract=4190623>.
32. De Ranieri E, McGillivray B, Swaminathan S, Samarasinghe M, Gruen L; 2017. Peer Review Congress.

33. Obrecht M, Tibelius K, D'Aloisio G. Examining the value added by committee discussion in the review of applications for research awards. *Research Evaluation*. 2007; 16(2):79–91. <https://doi.org/10.3152/095820207X223785>
34. Fogelholm M, Leppinen S, Auvinen A, Raitanen J, Nuutinen A, Väänänen K. Panel discussion does not improve reliability of peer review for medical research grant proposals. *Journal of clinical epidemiology*. 2012; 65(1):47–52. <https://doi.org/10.1016/j.jclinepi.2011.05.001> PMID: 21831594
35. Lawrence N, Cortes C. The NIPS Experiment; 2014. <http://inverseprobability.com/2014/12/16/the-nips-experiment>.
36. Pier E, Raclaw J, Kaatz A, Brauer M, Carnes M, Nathan M, et al. Your comments are meaner than your score: score calibration talk influences intra-and inter-panel variability during scientific grant peer review. *Research Evaluation*. 2017; 26(1):1–14. <https://doi.org/10.1093/reseval/rvw025> PMID: 28458466
37. Beygelzimer A, Dauphin Y, Liang P, Wortman Vaughan J. The NeurIPS 2021 Consistency Experiment; 2021. <https://blog.neurips.cc/2021/12/08/the-neurips-2021-consistency-experiment/>.
38. Rastogi C, Stelmakh I, Beygelzimer A, Dauphin YN, Liang P, Vaughan JW, et al. How do Authors' Perceptions of their Papers Compare with Co-authors' Perceptions and Peer-review Decisions? arXiv:2211.12966. 2022;.
39. Fortnow L. Viewpoint time for computer science to grow up. *Communications of the ACM*. 2009; 52(8):33–35. <https://doi.org/10.1145/1536616.1536631>
40. Parhami B. Low acceptance rates of conference papers considered harmful. *Computer*. 2016; 49(4):70–73. <https://doi.org/10.1109/MC.2016.106>
41. Vardi MY. Divination by program committee. *Communications of the ACM*. 2017; 60(9):7–7.
42. Shah NB. An Overview of Challenges, Experiments, and Computational Solutions in Peer Review (Extended Version); 2022. <http://bit.ly/PeerReviewOverview> (Abridged version published in the *Communications of the ACM*).
43. Teplitskiy M, Ranu H, Gray GS, Menietti M, Guinan E, Lakhani KR. Do Experts Listen to Other Experts?: Field Experimental Evidence from Scientific Peer Review; 2019.