

On the Privacy-Utility Tradeoff in Peer-Review Data Analysis

Wenxin Ding¹, Nihar B. Shah^{1, 2}, Weina Wang¹

¹ Computer Science Department, ² Machine Learning Department
Carnegie Mellon University

{wenxind@andrew, nihars@cs, weinaw@cs}.cmu.edu

Abstract

1 A major impediment to research on improving peer review
2 is the unavailability of peer-review data, since any release of
3 such data must grapple with the sensitivity of the peer review
4 data in terms of protecting identities of reviewers from au-
5 thors. We posit the need to develop techniques to release peer-
6 review data in a privacy-preserving manner. Identifying this
7 problem, in this paper we propose a framework for privacy-
8 preserving release of certain conference peer-review data —
9 distributions of ratings, miscalibration, and subjectivity —
10 with an emphasis on the accuracy (or utility) of the released
11 data. The crux of the framework lies in recognizing that a
12 part of the data pertaining to the reviews is already avail-
13 able in public, and we use this information to post-process
14 the data released by any privacy mechanism in a manner that
15 improves the accuracy (utility) of the data while retaining the
16 privacy guarantees. Our framework works with any privacy-
17 preserving mechanism that operates via releasing perturbed
18 data. We present several positive and negative theoretical re-
19 sults, including a polynomial-time algorithm for improving
20 on the privacy-utility tradeoff.

1 Introduction

21 A fair and efficient peer-review process is of utmost impor-
22 tance to the development of scientific research. There are,
23 however, a large number of challenges in peer review, per-
24 taining to its fairness and efficiency. Consequently there is
25 an overwhelming desire to “fix” the “broken” peer review
26 process (Rennie 2016; McCook 2006). And taking heed to
27 this call, there is a growing amount of research on this topic.

28 Research on improving peer review suffers from a con-
29 siderable handicap – unavailability of data (Balietti, Gold-
30 stone, and Helbing 2016; Tomkins, Zhang, and Heavlin
31 2017; Squazzoni et al. 2020; Schroter, Loder, and Godlee
32 2020). Concealing the identities of reviewers from authors of
33 any paper is paramount in most peer review systems. Thus
34 releasing any peer review data is fraught with the risk of
35 compromising on this privacy. As noted by Balietti (2016):

36
37 *“The main reason behind the lack of empirical stud-
38 ies on peer-review is the difficulty in accessing data.
39 In fact, peer-review data is considered very sensitive,*

*and it is very seldom released for scrutiny, even in an
anonymous form.”* 40 41

42 Although there is a large body of research on the topic
43 of privacy in various domains, not much privacy research
44 directly targets the application of peer review. In an influ-
45 ential recent paper about peer review, Tomkins, Zhang, and
46 Heavlin (2017) highlight the challenges they faced in this
47 respect and their consequent inability to release data:

48 *“We would prefer to make available the raw data used
49 in our study, but after some effort we have not been able
50 to devise an anonymization scheme that will simultane-
51 ously protect the identities of the parties involved and
52 allow accurate aggregate statistical analysis. We are
53 familiar with the literature around privacy preserving
54 dissemination of data for statistical analysis and feel
55 that releasing our data is not possible using current
56 state-of-the-art techniques.”*

57 We thus posit the need to develop techniques to help re-
58 lease peer-review data while ensuring that identities of re-
59 viewers of any paper are protected. With this motivation, we
60 focus on the privacy-utility tradeoff in releasing certain con-
61 ference peer-review data. The data to be released comprises
62 distributions of the ratings or miscalibration or subjectivity
63 in the peer-review process. The notion of privacy we con-
64 sider is quite general – our techniques apply to any notion of
65 privacy which operates by perturbing the data, including dif-
66 ferential privacy. We design a framework to improve in this
67 tradeoff by improving the utility (accuracy) while retaining
68 privacy guarantees.

69 Our work relies on the key observation that a non-trivial
70 part of conference peer-review data is already available in
71 the public domain. We design techniques which use this pub-
72 licly available information to post-process the data released
73 by any privacy mechanism. Our approach is guided by the
74 following four desiderata for such a post processing:

- D1 Under no circumstances should the accuracy decrease
after applying the algorithm. 75 76
- D2 Under no circumstances should the privacy guarantee be
compromised after applying the algorithm. 77 78
- D3 The algorithm should have a computational complexity
polynomial in the number of reviewers and papers.¹ 79 80

¹In typical conferences, the number of papers per reviewer and

D4 In special cases where an exact answer can be easily obtained from public data, the algorithm should also return the same answer with no error. (This is defined formally in Section 4.3.)

Our technical contributions (detailed in Section 4) towards this problem are as follows. We use the straightforward observation that projecting the (noisy) output of the privacy mechanism on the convex hull of all possible true values is desirable from the perspective of the desiderata. We prove that, however, such a projection is NP-hard (via reducing the ℓ -partition problem). We then design a polynomial-time computable algorithm which projects the noisy output of the privacy mechanism on a convex set containing all possible true values, and satisfies the four desiderata listed above. As a result of independent interest, we also prove that the more obvious approach of projecting on the set of all true values (instead of a convex set containing them) can, in fact, reduce the accuracy.

Finally, in Appendix C, we conduct synthetic simulations, which reveal that our methods can yield considerable improvements in the privacy-utility tradeoff as compared to standard approaches. The associated code for our algorithm is available at <https://github.com/wenixind/privacy-utility-tradeoff-in-peer-review-data>.

2 Related Work

This work falls in the intersection of two lines of research: peer review and privacy.

Peer review: Peer review is the backbone of scientific research. There is an overwhelming desire in many domains of science and engineering for improving peer review, and consequently, there are many past works on the topic of either evaluating the efficacy of peer review or improving the peer review process (Peters and Ceci 1982; Kliwer et al. 2004; Bennett, Jagsi, and Zietman 2018; Mavrogenis, Quaille, and Scarlet 2020; Bernard 2018; Snodgrass 2006; Scott 1974; Lindsey 1988; Douceur 2009; Reinhart 2009). These works, however, largely focus on the journal reviewing setup that is common in non-computer-science fields, whereas our focus is on the conference reviewing setting which is more common in computer science.

The number of submissions to many computer science conferences, particularly to machine learning or artificial intelligence conferences, is growing near-exponentially and is presently in the several thousands. This rapid growth has spurred a considerable amount of recent research on peer review in computer science. These works include those on handling problems related to reviewer-assignment (Goldsmith and Sloan 2007; Charlin and Zemel 2013; Welch 2014; Stelmakh, Shah, and Singh 2018; Kobren, Saha, and McCallum 2019), miscalibration (Roos, Rothe, and Scheuermann 2011; Ge, Welling, and Ghahramani 2013; Wang and Shah 2019), subjectivity (Noothigattu, Shah, and Procaccia 2018), biases (Tomkins, Zhang, and Heavlin 2017; Stelmakh, Shah, and Singh 2019), strategic behavior (Baliatti, Goldstone, and Helbing 2016; Xu et al. 2019;

the number of reviewers per paper are both constants (Shah et al. 2018).

Stelmakh, Shah, and Singh 2020) and others (Cabanac and Preuss 2013; Fiez, Shah, and Ratliff 2019; Lawrence and Cortes 2014; Shah et al. 2018; Stelmakh et al. 2020). In particular, as will be detailed later, our work is also useful towards releasing data pertaining to miscalibration and subjectivity, thereby helping in the understanding and mitigation of these problems.

Privacy: Privacy-preserving data analytics has been receiving rapidly increasing attention as the big-data regime emerges. There is a large body of research that investigates formal notion of privacy and quantifies the tradeoff between privacy and utility (see, e.g., Dwork et al. 2006b; Dwork 2006; Blum, Ligett, and Roth 2008; Gaboardi et al. 2014; Wang, Ying, and Zhang 2016; Bun, Ullman, and Vadhan 2018). Among these studies, differential privacy (Dwork et al. 2006b; Dwork 2006) has become the de facto standard and has been applied to many areas.

In this paper, we investigate the privacy-utility tradeoff for publishing histograms of peer-review data. Privacy-preserving release of histograms has been a major focus of the literature (Chawla et al. 2005; Dwork et al. 2006a; Hay et al. 2010; Li et al. 2010; Bassily and Smith 2015; Balcer and Vadhan 2019; Abowd et al. 2019). To the best of our knowledge, existing techniques for improving the privacy-utility tradeoff are generally inadequate for the application of peer review since they do not take into account the special structures in peer-review data.

In particular, one special feature of peer-review data is that some specific part of the data such as scores received by papers is already published in its original, non-privacy-preserving form. This provides us an opportunity to utilize the “consistency” with public knowledge. The concept of consistency, with different problem-specific meanings, has been investigated by existing work for privacy-preserving algorithms. Hay et al. 2010 improves accuracy by assuring consistency among answers to different queries. The closest work to ours is the privacy-preserving approach for US Census (Abowd et al. 2019), where consistency with public data is a requirement and it is in the form of a set of linear constraints. In contrast, we are not subject to a strict requirement of consistency, but instead, we exploit consistency as a method to improve accuracy. Additionally, the consistency with public knowledge in our problem is of a more combinatorial nature. As we discuss in the sequel, the idiosyncratic nature of the peer-review setting implies that one can design methods tailored to this application which yield a (considerable) improvement in the privacy-utility tradeoff as compared to standard privacy mechanisms.

Peer review and privacy: An exception is the concurrent work (Jecmen et al. 2020) which considers releasing the reviewer-paper similarity matrix and source code for the reviewer assignment (whereas in contrast we consider releasing a function of the scores given by reviewers to papers). Their approach involves modifying and randomizing the reviewer-paper assignment process and their guarantees pertain to plausible deniability (that is, any reviewer may be assigned to any paper with a probability at most a certain value). On the other hand, we do not modify the peer-review process in any way, and instead use any privacy-preserving

195 data-release mechanism coupled with post processing of the
 196 data from peer review.

197 3 Background and problem setting

198 In this section, we provide some background on the peer
 199 review setting and privacy, and describe our problem setting
 200 in more detail.

201 3.1 Peer review

202 We consider a conference peer review setting, where there
 203 are n reviewers and m papers. We index the papers as $[m] =$
 204 $\{1, 2, \dots, m\}$ and the reviewers as $[n] = \{1, 2, \dots, n\}$.²
 205 For simplicity we assume that the number of papers re-
 206 viewed by each reviewer is the same for all reviewers –
 207 denoted as ℓ , and that the number of reviewers reviewing
 208 each paper is the same for all papers – denoted as k .³ Conse-
 209 quently, we have the relation $n\ell = mk$. All four parameters
 210 (n, m, ℓ, k) are public knowledge.

211 Each review comprises a real-valued score. We assume
 212 that all papers and all associated reviews (that is, the set of
 213 scores received by each paper) are public knowledge (e.g., in
 214 conferences such as ICLR and others on the OpenReview.net
 215 review platform). The list of all reviewers is also available
 216 publicly (such a list is released by many conferences). How-
 217 ever, importantly, the identity of which reviewer reviewed
 218 which paper is private.

219 We now introduce notation to describe the score given
 220 in any review. If reviewer $j \in [n]$ reviews paper $i \in [m]$,
 221 then we use $s_{ij} \in \mathbb{R}$ to denote the score of this review.
 222 This score is private in the sense that the identity of the re-
 223 viewer who gives this score is not publicly available. How-
 224 ever, for each paper $i \in [m]$, the multiset $\{s_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i\}$ is public.

225 This setting can be described by a bipartite graph, as
 226 shown in Figure 1. The bipartite graph has two disjoint sets
 227 of vertices, $[m]$ and $[n]$ representing the sets of papers and
 228 reviewers, respectively. In private data (Figure 1a), an edge
 229 exists between any vertex (paper) $i \in [m]$ and any ver-
 230 tex (reviewer) $j \in [n]$ if reviewer $j \in [n]$ reviews paper
 231 $i \in [m]$. We associate each edge (i, j) with the score s_{ij} .
 232 The edges (and their associated scores) are all private. The
 233 private data is accessible to the program chairs of the confer-
 234 ence. In public data (Figure 1b), for each vertex (paper) in
 235 $[m]$, the weights of the edges connected to it are known pub-
 236 licly. However, the edges of the graph are not known. Note
 237 that in both public and private data, identities of papers and
 238 reviewers are known.

239 There are various quantities of interest for release that
 240 we consider in this work. An intermediate set of terms to-
 241 wards these quantities is the multiset $\{w_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i \in [m]\}$ discussed below, which we refer to as the set of weights. This multiset can be computed from the scores $\{s_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i \in [m]\}$.

²We follow the standard convention of using $[\beta]$ to represent the set $\{1, 2, \dots, \beta\}$ for any positive integer β .

³Our work is also applicable to the most general setting in which different reviewers and/or different papers have different loads. We discuss this in Section 4.3.

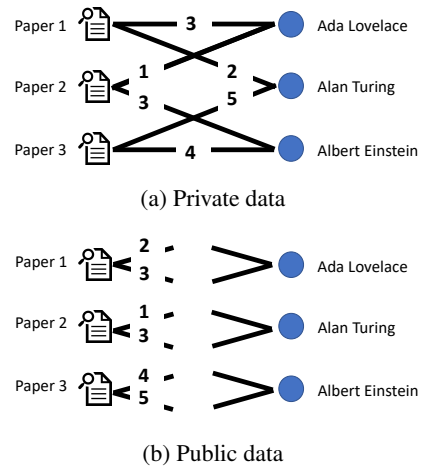


Figure 1: An illustration of the data (a) available privately to the program chairs of the conference, and (b) available to the public under increasingly popular ‘open review’ paradigms in computer science.

246 We now discuss three such choices of $\{w_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i \in [m]\}$, and subsequently describe the data we aim to release.

- 247 • **Reviewer ratings.** In this case, the mapping from scores to weights is simply the identity mapping: 248 249 250

$$w_{ij} = s_{ij}. \quad (3.1)$$

- 251 • **Miscalibration.** Miscalibration pertains to the problem of 252 strictness or leniency of different reviewers, which can 253 govern the fate of papers (Roos, Rothe, and Scheuermann 2011; Ge, Welling, and Ghahramani 2013; Wang and Shah 2019). In order to understand the amount of 254 miscalibration, it is instructive to see the difference be- 255 tween the scores given by a reviewer and the scores given 256 by other reviewers for the same papers. To this end, we 257 let w_{ij} denote the miscalibration in any individual review 258 (for any paper i by any reviewer j): 259 260

$$w_{ij} = s_{ij} - \frac{1}{k-1} \sum_{j' \neq j} s_{ij'}. \quad (3.2)$$

- 261 • **Subjectivity.** Subjectivity is the problem that different re- 262 viewers place different emphasis on the various criteria 263 when making an overall decision for a paper (Lee 2015). 264 Techniques such as that proposed in Noothigattu, Shah, 265 and Procaccia 2018 can be used to normalize each score 266 in a manner that mitigates the subjectivity. Specifically, 267 the technique in Noothigattu, Shah, and Procaccia 2018 268 uses the public data to transform the score s_{ij} associated 269 to each review into a normalized version, say, \tilde{s}_{ij} . We can 270 then set $w_{ij} = \tilde{s}_{ij}$ for every review, and the algorithm 271 in this paper will help release statistics of these normal- 272 ized scores. A second use case we consider is to better 273 understand and investigate the issue of subjectivity, by 274 releasing the amount of subjectivity present in the sys- 275 tem, that is, the aggregate difference of reviewers’ scores

and their normalized scores. Concretely in this case, after obtaining the normalized scores $\{\tilde{s}_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i \in [m]\}$, we set $w_{ij} = s_{ij} - \tilde{s}_{ij}$ for every review.

Analogous to the scores s_{ij} 's, the weights w_{ij} 's are also associated to public and private components. We can use the same bipartite graphs as in Figure 1 to represent the setting with weights. In particular, the private data continues to include the edges pertaining to which reviewer reviewed which paper. The private data also includes the weight w_{ij} on each edge (i, j) representing the weight that reviewer $i \in [n]$ gives to paper $j \in [m]$. The private data is depicted in Figure 1a where in this interpretation, the values on the edges represent the weights. The public data, as in the case of scores, only includes the multiset of weights received by any paper, that is, the public data comprises the multisets $\{w_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i\}$ for every paper $i \in [m]$. The public data is depicted in Figure 1b where the values on the edges represent the weights.

It is very important to note the following two properties in the transformation of scores s_{ij} to weights w_{ij} for each of the aforementioned choices. First, clearly, given access to all private scores, all weights can be computed. Second, the public weights (that is, the multisets $\{w_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i\}$ for every paper $i \in [m]$) can be computed using only the publicly available score data (that is, the multisets $\{s_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i\}$ for every paper $i \in [m]$). This relation between the public (or private) weights and public (or private) scores allows us to interchange them in the graphs in Figure 1.

For each reviewer $j \in [n]$, let \mathcal{Y}_j be the set of all papers reviewed by reviewer j , that is, $\mathcal{Y}_j = \{i \in [m] \mid \text{reviewer } j \text{ reviews paper } i\}$. Let y_j denote the mean weight of reviewer j :

$$y_j = \frac{1}{\ell} \sum_{i \in \mathcal{Y}_j} w_{ij}. \quad (3.3)$$

Note that since the identity of the reviewer in any review is private, the values of \mathcal{Y}_j and y_j in general cannot be computed from the public data.

Quantity to be released: The quantity of interest is the histogram of the mean weights per reviewer, represented by the *sorted* version of the mean-weight vector (y_1, y_2, \dots, y_n) , which we denote by $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_n^*)$. Then $\theta_1^* \leq \dots \leq \theta_n^*$ and the multiset $(\theta_1^*, \theta_2^*, \dots, \theta_n^*)$ equals the multiset (y_1, y_2, \dots, y_n) . We call θ^* the true sorted mean-weight vector. According to the applications discussed above, the vector θ^* can either represent the mean scores per reviewer or capture the amount of miscalibration, or subjectivity in the reviews.

Our goal is to release θ^* , while ensuring privacy of reviewer identities. When the underlying weights are equal to the scores, the sorted mean-weight vector (that is, the histogram of scores) is commonly released by various conferences (Shah et al. 2018). These are however usually released without any privacy considerations, and our work addresses privacy-preserving release with high accuracy. Addressing the issues of miscalibration and subjectivity is extremely important for fair and high-quality peer review (Ge, Welling,

and Ghahramani 2013; Wang and Shah 2019; Roos, Rothe, and Scheuermann 2011; Lee 2015; Noothigattu, Shah, and Procaccia 2018; Siegelman 1991; Kerr, Tolliver, and Petree 1977), and releasing the statistics pertaining to the amount of miscalibration or subjectivity can considerably help both research and policy-design regarding these issues.

Publishing histograms of datasets in a privacy-preserving manner has been a central objective in the literature of privacy research (Chawla et al. 2005; Dwork et al. 2006a; Hay et al. 2010; Li et al. 2010; Bassily and Smith 2015; Balcer and Vadhan 2019). Histograms are typically the most frequently used statistics in official reports, and more importantly, they form the basis for more complicated statistical analysis. To the best of our knowledge, existing techniques for improving the privacy-utility tradeoff (for histograms or otherwise) are generally inadequate for the application of peer review since they do not take into account the special structures in peer-review data. Our goal is to use the specific type of publicly available data in peer review in order to improve the privacy-utility tradeoff.

Finally we note that the high-level ideas behind our proposed algorithm are more general and may also be used to improve the utility of the released data for settings beyond the sorted mean-weight vector. We revisit this point later in the paper.

3.2 Privacy

To protect the privacy of reviewers, we consider privacy-preserving mechanisms that (randomly) perturb the quantities of interest. By virtue of the random perturbation, the privacy mechanism makes it hard to infer each individual reviewer's scores given to papers from the noisy data. Specifically, we consider any privacy mechanism that releases a vector $\mathbf{r} = (r_1, r_2, \dots, r_n) \in \mathbb{R}^n$ obtained by randomly perturbing the sorted mean-weight vector $(\theta_1^*, \theta_2^*, \dots, \theta_n^*) \in \mathbb{R}^n$. An example of a privacy mechanism is the Laplace mechanism, which satisfies the popular notion of differential privacy (Dwork et al. 2016; Dwork and Roth 2014) in which $(r_1, r_2, \dots, r_n) = (\theta_1^*, \theta_2^*, \dots, \theta_n^*) + (\eta_1, \eta_2, \dots, \eta_n)$. Here $\eta_1, \eta_2, \dots, \eta_n$ are i.i.d. random variables drawn from a zero-mean Laplace distribution.

3.3 Utility (Accuracy)

Let $\mathbf{t} = (t_1, \dots, t_n)$ denote the final output (after post-processing) that is released. We measure the utility or accuracy of the output in terms of its *mean squared error* with respect to the true value of the vector θ^* , that is, $\mathbb{E} \left[\sum_{i=1}^n (\theta_i^* - t_i)^2 \right]$. We say that an (possibly random) output $\mathbf{t} = (t_1, \dots, t_n)$ is more accurate than another output $\mathbf{t}' = (t'_1, \dots, t'_n)$ with respect to θ^* if

$$\mathbb{E} \left[\sum_{i=1}^n (\theta_i^* - t_i)^2 \right] < \mathbb{E} \left[\sum_{i=1}^n (\theta_i^* - t'_i)^2 \right]. \quad (3.4)$$

3.4 Goal

Our goal is to design algorithms to process the data output by the privacy-preserving mechanism, \mathbf{r} , before its actual re-

383 lease in a manner that improves the privacy-utility tradeoff.
 384 Specifically, we aim to satisfy the four desiderata D1–D4
 385 listed in Section 1.

386 4 Main Results

387 4.1 Approach

388 We first derive a representation of the set of all possible val-
 389 ues in the sorted mean-weight vector θ^* based on the public
 390 data. For any paper $i \in [m]$, we use $x_{i1}, x_{i2}, \dots, x_{ik}$ to de-
 391 note the k weights on edges connected to that vertex (paper)
 392 in the public data, listed in an arbitrary order. Note that the
 393 second subscript of x_{ij} does not correspond to a reviewer
 394 identity. The multiset $\{x_{i1}, \dots, x_{ik}\}$ is available publicly,
 395 and for each $i \in [m]$, the multiset $\{x_{i1}, \dots, x_{ik}\}$ is identi-
 396 cal to the multiset $\{w_{ij} \mid \text{reviewer } j \in [n] \text{ reviews paper } i\}$.
 397 Let \mathcal{G} be a set of weighted bipartite graphs comprising all
 398 valid reviewer-paper weights based on the public data, that
 399 is, each member of \mathcal{G} satisfies:

- 400 • It is a bipartite graph, with the vertices in the two parts
 401 corresponding to papers $[m]$ and reviewers $[n]$.
- 402 • All vertices in $[m]$ are k regular and all vertices in $[n]$ are
 403 ℓ regular.
- 404 • The k edges incident on any vertex $i \in [m]$ have weights
 405 $x_{i1}, x_{i2}, \dots, x_{ik}$.

406 Furthermore, for any graph $g \in \mathcal{G}$, any paper i , and any
 407 reviewer j , we define $w_{ij}(g)$ equal to the weight of edge
 408 between i and j if this edge exists, and $w_{ij}(g) = 0$ other-
 409 wise. Then the set Θ , that comprises all possible values of
 410 the sorted mean-weight vector based on public data, is given
 411 by

$$412 \Theta = \left\{ \theta \in \mathbb{R}^n \mid \theta_1 \leq \dots \leq \theta_n, \exists g \in \mathcal{G} \text{ such that} \right. \\
 413 \left. \theta_j = \frac{1}{\ell} \sum_{i=1}^m w_{ij}(g) \text{ for all } j \in [n] \right\}. \quad (4.1)$$

412 Note that the true paper-reviewer graph is also a member of
 413 \mathcal{G} and the true sorted weight vector $\theta^* \in \Theta$. Throughout
 414 this section, we consider algorithms based on projecting the
 415 noisy data on certain sets. To this end, for any set $\mathcal{C} \subseteq \mathbb{R}^n$,
 416 we define the projection of vector \mathbf{r} on the set \mathcal{C} as

$$417 \operatorname{argmin}_{\theta \in \mathcal{C}} \sum_{i=1}^n (\theta_i - r_i)^2. \quad (4.2)$$

418 When the privacy-preserving algorithm perturbs the true
 419 sorted mean-weight vector θ^* , the resulting noisy mean-
 420 weight vector \mathbf{r} may not lie in the set Θ . It is thus intuitive
 421 to instead replace the resulting vector with the vector in Θ
 422 closest to it, that is, to instead output the projection (4.2) of
 423 the vector \mathbf{r} with the choice $\mathcal{C} = \Theta$.

424 The following result shows that this intuitive approach can
 425 actually *increase* the error. We state and prove this result
 426 concretely in the case of additive Laplace noise, but as seen
 427 in the proof, the result is much more general.

427 **Proposition 4.1.** *There exists a review setting such that the*
 428 *true sorted mean-weight vector θ^* , the noisy mean-weight*
 429 *vector \mathbf{r} obtained by adding Laplace noise with zero mean*

and a fixed, non-zero variance to θ^* , and the output \mathbf{t} of the
 projection (4.2) of \mathbf{r} on the set $\mathcal{C} = \Theta$, are related as

$$430 \mathbb{E} \left[\sum_{i=1}^n (t_i - \theta_i^*)^2 \right] > \mathbb{E} \left[\sum_{i=1}^n (r_i - \theta_i^*)^2 \right], \quad (4.3)$$

431 where the expectation is taken with respect to the noise dis-
 432 tribution.

433 The proposition implies that using the closest valid vector
 434 violates desideratum D1 of not reducing the accuracy. The
 435 proof of Proposition 4.1 is given in Appendix D.1.

436 Consequently, in order to ensure desideratum D1 of not
 437 reducing the accuracy is met, we project the noisy data onto
 438 a convex set that contains Θ . The following proposition
 439 (proved in Appendix D.2) indicates that such projection can
 440 never harm the accuracy, and is a straightforward application
 441 of the fact that projection on to convex sets is non-expansive.
 442 Note that projection methods have been used in the litera-
 443 ture (Hay et al. 2010; Abowd et al. 2019) and it is known
 444 that projection does not decrease accuracy. For complete-
 445 ness, we include Proposition 4.2 here as the specific form of
 446 such results for our problem.

447 **Proposition 4.2.** *Consider any true sorted mean-weight*
 448 *vector θ^* , and any arbitrary (noisy mean-weight) vector*
 449 *\mathbf{r} . Let \mathcal{C} be any closed convex set such that $\Theta \subseteq \mathcal{C}$. Let*
 450 *$\mathbf{t} = (t_1, \dots, t_n)$ be the projection of \mathbf{r} on to set \mathcal{C} as in (4.2).*
 451 *Then it must be that*

$$452 \sum_{i=1}^n (t_i - \theta_i^*)^2 \leq \sum_{i=1}^n (r_i - \theta_i^*)^2. \quad (4.4)$$

453 Since proposition 4.2 holds for all \mathbf{r} , it follows that if \mathbf{r} is
 454 obtained by randomly perturbing θ^* , then

$$455 \mathbb{E} \left[\sum_{i=1}^n (t_i - \theta_i^*)^2 \right] \leq \mathbb{E} \left[\sum_{i=1}^n (r_i - \theta_i^*)^2 \right]. \quad (4.5)$$

456 Moreover, for two closed convex sets $\mathcal{C}_1 \subseteq \mathcal{C}_2$ both contain-
 457 ing Θ , if we have a projection on \mathcal{C}_2 , then further projecting
 458 it on \mathcal{C}_1 can never increase the error and can possibly de-
 459 crease the error. **Our goal thus is to project the noisy data**
 460 **on to a (small) convex set that contains all possible true**
 461 **values.**

462 4.2 NP-hardness of Projection onto Convex Hull

463 The smallest convex set that contains Θ is the convex hull
 464 of Θ . Observe that if we could project on to the convex hull,
 465 then it can also be used to improve upon the projection on
 466 any other convex set. Specifically, if \mathbf{t} is the projection of
 467 the perturbed data \mathbf{r} on some convex set that contains Θ ,
 468 and if \mathbf{t}' is the projection of \mathbf{t} on convex-hull(Θ), then with
 an argument identical to that in Proposition 4.2 we have that

$$469 \sum_{i=1}^n (t'_i - \theta_i^*)^2 \leq \sum_{i=1}^n (t_i - \theta_i^*)^2.$$

470 Consequently, in this section we consider the goal of pro-
 471 jecting the noisy data onto the convex hull of Θ . In this case,
 472 the final result we will output can be represented as choosing
 473 $\mathcal{C} = \text{convex-hull}(\Theta)$ in Equation (4.2). Unfortunately, as we
 474 show below, projection onto convex-hull(Θ) is NP-hard.

475 **Theorem 4.3.** When $k = \ell > 2$, $m = n$ and n is a multi-
 476 ple of ℓ , the problem of projecting noisy data onto convex-
 477 hull(Θ) is NP-hard.

478 We prove this result via reducing the ℓ -Partition prob-
 479 lem to the projection problem. Given any instance of an ℓ -
 480 Partition problem, which is a multiset of integers, we can
 481 construct a conference where each paper has a weight from
 482 the multiset. We can construct a vector such that the projec-
 483 tion result can directly answer the ℓ -Partition problem. The
 484 complete proof of Theorem 4.3 is provided in Appendix D.3.

485 4.3 An Efficient Algorithm

486 In this section, we present an algorithm that meets the four
 487 desiderata D1–D4 listed in Section 1. Since we cannot effi-
 488 ciently project on to the convex hull of Θ , we must make do
 489 with a larger convex set that contains Θ . We use desidera-
 490 tum D4 for guidance on what constitutes a reasonably small
 491 set and associated projection.

492 **Axioms defining desideratum D4** Recall that desidera-
 493 tum D4 says that the algorithm should automatically recover
 494 the ground truth when the structure of the public data is sim-
 495 ple enough. More concretely, we benchmark any algorithm
 496 using the following axiomatic properties:

- 497 A1 When all weights are identical, the projection should re-
 498 sult in a vector whose entries are all the same as the
 499 weight. Formally, if $x_{ij} = z \forall i \in [m], j \in [k]$ for
 500 some z , then the output \mathbf{t} of the algorithm must be
 501 $t_1 = t_2 = \dots = t_n = z$.
- 502 A2 When $\ell = 1$ (that is, each reviewer reviews 1 pa-
 503 per), the projection of any noisy data should result in a
 504 sorted vector of all weights. Formally, if $\ell = 1$ then
 505 the output \mathbf{t} of the algorithm must be $(t_1, t_2, \dots, t_n) =$
 506 $\text{sorted}(x_{11}, x_{21}, \dots, x_{n1})$.
- 507 A3 When all but one papers have all zero weights, the
 508 projection of any noisy data should result in a sorted
 509 vector with $(n - k)$ zero entries and the remain-
 510 ing entries equal to $\frac{1}{\ell}$ of the weights for the pa-
 511 per that does not have all-zero weights. Formally, if
 512 $x_{ij} = 0 \forall i \in \{2, \dots, m\}, j \in [k]$, then the out-
 513 put \mathbf{t} of the algorithm must be $(t_1, t_2, \dots, t_n) =$
 514 $\text{sorted}(\frac{x_{11}}{\ell}, \frac{x_{12}}{\ell}, \dots, \frac{x_{1k}}{\ell}, 0, \dots, 0)$.

515 **High-level idea behind the algorithm** The main idea be-
 516 hind our algorithm comprises the following three steps:

- 517 I. From the public data, take all tuples of size ℓ contain-
 518 ing weights from different papers into consideration.
- 519 II. Use them to construct lower and upper bounds on ev-
 520 ery entry of (the unknown vector) θ^* .
- 521 III. Project the noisy data \mathbf{r} on the set specified by the
 522 aforementioned lower and upper bounds, along with
 523 any other problem-specific (convex) constraints.

524 As one can intuitively see, these three steps imply a pro-
 525 jection of the noisy data on a convex set which includes all
 526 valid values of the true data, and hence from Proposition 4.2
 527 it will not reduce the utility (desideratum D1). Moreover, the
 528 entire algorithm uses only the public data along with the vec-
 529 tor \mathbf{r} released by the privacy mechanism, and hence does not
 530 compromise privacy (desideratum D2). The idea is general

enough to be applied to many forms of the noisy data, and
 in what follows, we apply it to release the histogram of the
 true sorted mean-weight vector. Of course, the devil lies in
 the details of how these steps are executed, which will deter-
 mine whether the designed algorithm meets desiderata D3
 and D4.

Full algorithm description We now describe our algo-
 rithm in full detail. (We also provide an illustrative exam-
 ple in Appendix B.) Recall that we use $x_{i1}, x_{i2}, \dots, x_{ik}$
 to represent the k weights on edges connected to any vertex
 (paper) $i \in [m]$ in the public data. Note that since reviewer
 identities are not available publicly, the second subscript “ j ”
 in “ x_{ij} ” has no particular meaning other than capturing the
 fact that each paper has k weights. We use matrix X to dis-
 play all the weights in the public data where row i column j
 of X has value x_{ij} .

I. Valid weight tuples We define a weight tuple as a mul-
 tiset of ℓ real values. We say that a tuple is a *valid weight*
tuple if it consists of ℓ weights from distinct papers. In other
 words, a valid weight tuple contains ℓ entries of matrix X
 where no two entries are from the same row in X . We com-
 pute Ω' as the list of all valid weight tuples. In other words,
 Ω' contains all the possible weight tuples given by a re-
 viewer. We sort the list Ω' based on the mean weight of the
 weight tuples (breaking ties arbitrarily), and henceforth use
 the notation Ω for this sorted list.

II. Lower and upper bounds We now compute lower and
 upper bounds on each entry of θ^* based only on the public
 data. We create a graph G which has all weight tuples in Ω
 as its vertices. Since each weight tuple in Ω corresponds to a
 vertex in graph G , we use the terms “weight tuples in Ω ” and
 “vertices in G ” interchangeably. There is an edge between
 two vertices if the two weight tuples do not contain weights
 corresponding to the same entry in X . Then for each vertex,
 we define its left chain and right chain as follows. Recall
 that Ω is a sorted list and all of its entries, which are weight
 tuples, are totally ordered. We use the indices of the tuples
 (vertices) in this ordering for the following definitions.

Definition 4.4. For any vertex ν in G , a left chain (resp.
 right chain) of ν is a simple path in G from ν to another ver-
 tex such that the indices of the vertices in this path decrease
 (resp. increase) starting from ν .

We also define the *length of a chain* to be the number of
 vertices in the chain. For each vertex, we compute the length
 of its longest left chain and the length of its longest right
 chain using dynamic programming. To compute the length
 of the longest left chain of a vertex ν in G , we check the
 length of the longest left chain of all its neighbors at lower
 indices. Then the length of the longest left chain of ν is
 the maximum of these neighbors’ longest left chain lengths
 plus one. Similarly, to compute the length of the longest
 right chain of ν , we check the length of the longest right
 chain of all its neighbors at higher indices. The length of the
 longest right chain of ν is the maximum of its neighbors’
 longest right chain lengths plus one. We store the length of
 the longest left and right chain of each vertex for subsequent
 use in the algorithm.

The algorithm to compute a lower bound on θ_i^* for each

Algorithm 1: Computation of lower bounds

Input: matrix X of weights, sorted list of weight tuples Ω

Initialize $i = 1$, set $w \in \mathbb{R}^\ell$ as the first tuple in Ω , and all entries of X are unmarked.

while $i \leq n$ **do**

 for each weight in the tuple w , find its corresponding entry in matrix X and mark the entry

if length of tuple w 's longest left chain $\geq i$ **and** number of unmarked entries on each row of

$X \leq n - i$ **then**

 lower bound on θ_i^* = mean of all entries of tuple w
 $i += 1$

end if

 set w as the next tuple in Ω

end while

uniform reviewer and paper loads. First, if the papers are reviewed by different number of reviewers, the algorithm above continues to work. Now if the reviewers review different numbers of papers, then we make the following modification to the algorithm. Let $\mathcal{L} \subset [m]$ denote the set of all paper loads on the reviewers, that is, $\ell \in \mathcal{L} \iff$ there is a reviewer who reviews exactly ℓ papers. Then the set Ω' computed in the first step of the algorithm includes all weight tuples of size ℓ for every $\ell \in \mathcal{L}$. The remainder of the algorithm remains identical to that described above. The proof of correctness in these settings follows from the same arguments (given in Appendix D.5) as those for the setting of uniform reviewer and paper loads. The algorithm continues to have a computational complexity that is polynomial in n and m (where we continue to assume that the maximum reviewer and paper loads are constants (Shah et al. 2018)).

Guarantees of Our Algorithm In this section, we evaluate our algorithm with respect to the four desiderata listed in Section 1. We first prove the correctness of the algorithm in terms of projection on to an appropriate set.

Theorem 4.6. *The algorithm projects noisy data onto a convex set that contains all true values.*

The proof of this theorem is given in Appendix D.5. This result, combined with Theorem 4.2, guarantees that our algorithm does not increase the error. Thus, our algorithm satisfies desideratum D1. In addition, since our algorithm uses only the public data for post processing, it satisfies desideratum D2. We now discuss the computational complexity of our algorithm.

Theorem 4.7. *The algorithm has polynomial time complexity in the number of reviewers and the number of papers.*

Our algorithm thus satisfies desideratum D3. The proof of this theorem is given in Appendix D.6. To be clear, while the algorithm is polynomial time in n and m , it is exponential in ℓ . In practice ℓ is usually a small constant (Shah et al. 2018).

We finally visit desideratum D4 – of returning an exact answer when it can easily be deduced from public data.

Theorem 4.8. *The algorithm satisfies the axiomatic properties A1, A2 and A3 defined in Section 4.3.*

The proof of this theorem is given in Appendix D.7. We have thus shown that our proposed algorithm meets all four desiderata D1–D4.

5 Conclusion

We take the first steps towards designing methods for privacy-preserving release of peer-review data, and posit the need for much more research on this topic to address the important challenge of improving peer review. While we addressed a certain type of peer-review data, it is of theoretical and practical interest to enable privacy-preserving release of more peer-review data such as properties of the reviewer graph, reviewer bids, and other functions of the scores. Moreover, it is of interest to design methods that can utilize data from multiple conferences, while preserving the privacy in each conference, for improving the peer-review process in any subsequent conference.

589 $i \in [n]$ is presented in Algorithm 1. In more detail, the algorithm uses two criteria to determine if mean of a weight tuple is a lower bound on θ_i^* . The criteria are

590
591
592 C1 The longest left chain of the tuple has length at least i .
593 C2 In X , after we mark the ℓ weights from each tuple considered so far, each row has at most $n - i$ unmarked entries.

594
595
596 The intuition is as follows. We call the n weight tuples that compute θ^* the true weight tuples. The true weight tuple with mean θ_i^* must have $i - 1$ weight tuples with smaller or equal mean to θ_i^* . No two reviewers give the same weight so no two true weight tuples contain weights from the same entry in X . Thus, criterion C1 is a necessary condition for a weight tuple to be the true weight tuple that computes θ_i^* . In addition, there are $n - i$ entries in θ^* whose values are no smaller than θ_i^* . Since no two true weight tuples contain weights from the same paper, each paper can have at most $n - i$ unused weights. Therefore, each row of X cannot have more than $n - i$ unmarked entries. Thus, criterion C2 is necessary for all weights to be assigned among the reviewers. For each entry $i \in [n]$, we choose the valid weight tuple with the smallest mean that satisfies criteria C1 and C2 as the lower bound on θ_i^* . Hence, it is a valid lower bound.

602
603
604
605
606
607
608
609
610
611
612 The computation of the upper bounds is analogous to that of lower bounds, and is presented in Appendix A in detail.

613
614
615
616
617 **III. Projection** Let L_i denote the lower bound we compute on θ_i^* and U_i denote the upper bound we compute on θ_i^* in part II above. The final output of our algorithm is the solution to the following optimization problem:

$$\begin{aligned} \operatorname{argmin}_{t \in \mathbb{R}^n} \sum_{i=1}^n (\eta_i - t_i)^2 \text{ such that } L_i \leq t_i \leq U_i \forall i \in [n], \\ \sum_{i=1}^n t_i = \frac{1}{\ell} \sum_{i=1}^m \sum_{j=1}^k x_{ij}, t_1 \leq t_2 \leq \dots \leq t_n. \end{aligned} \quad (4.6)$$

618 This convex optimization problem with a quadratic objective and $2n$ linear constraints can be solved efficiently.

619
620 *Remark 4.5* (Extension to non-uniform paper and reviewer loads). Our algorithm easily extends to the setting of non-

References

- 677
- 678 Abowd, J.; Kifer, D.; Moran, B.; Ashmead, R.; Leclerc, P.;
679 Sexton, W.; Garfinkel, S.; and Machanavajjhala, A. 2019.
680 Census TopDown: Differentially Private Data, Incremental
681 Schemas, and Consistency with Public Knowledge. US Cen-
682 sus Bureau. [https://github.com/uscsensusbureau/census2020-](https://github.com/uscsensusbureau/census2020-das-e2e/blob/master/doc/20190711_0945_Consistency_for_Large_Scale_Differentially_Private_Histograms.pdf)
683 [das-e2e/blob/master/doc/20190711_0945_Consistency_for_](https://github.com/uscsensusbureau/census2020-das-e2e/blob/master/doc/20190711_0945_Consistency_for_Large_Scale_Differentially_Private_Histograms.pdf)
684 [Large_Scale_Differentially_Private_Histograms.pdf](https://github.com/uscsensusbureau/census2020-das-e2e/blob/master/doc/20190711_0945_Consistency_for_Large_Scale_Differentially_Private_Histograms.pdf).
- 685 Babel, L.; Kellerer, H.; and Kotov, V. 1998. The k-
686 partitioning problem. *Mathematical Methods of Operations*
687 *Research* 47(1): 59–82.
- 688 Balcer, V.; and Vadhan, S. 2019. Differential Privacy on
689 Finite Computers. *Journal of Privacy and Confidentiality*
690 9(2). doi:10.29012/jpc.679.
- 691 Baliotti, S. 2016. Science is suffering because of peer re-
692 view’s big problems. *The Conversation*.
- 693 Baliotti, S.; Goldstone, R. L.; and Helbing, D. 2016. Peer re-
694 view and competition in the Art Exhibition Game. *Proceed-*
695 *ings of the National Academy of Sciences* 113(30): 8414–
696 8419.
- 697 Bassily, R.; and Smith, A. 2015. Local, Private, Efficient
698 Protocols for Succinct Histograms. In *Proc. Ann. ACM*
699 *Symp. Theory of Computing (STOC)*, 127–135. Portland,
700 OR.
- 701 Bauschke, H. H.; Combettes, P. L.; et al. 2011. *Convex anal-*
702 *ysis and monotone operator theory in Hilbert spaces*, vol-
703 *ume 408*. Springer.
- 704 Bennett, K. E.; Jagsi, R.; and Zietman, A. 2018. Radiation
705 oncology authors and reviewers prefer double-blind peer re-
706 view. *Proceedings of the National Academy of Sciences*
707 115(9): E1940–E1940.
- 708 Bernard, C. 2018. Gender Bias in Publishing: Double-Blind
709 Reviewing as a Solution? *Eneuro* 5(3).
- 710 Blum, A.; Ligett, K.; and Roth, A. 2008. A learning the-
711 ory approach to non-interactive database privacy. In *Proc.*
712 *Ann. ACM Symp. Theory of Computing (STOC)*, 609–618.
713 Victoria, Canada.
- 714 Bun, M.; Ullman, J.; and Vadhan, S. 2018. Fingerprint-
715 ing Codes and the Price of Approximate Differential Pri-
716 vacy. *SIAM Journal on Computing* 47(5): 1888–1938.
717 doi:10.1137/15M1033587. URL [https://doi.org/10.1137/](https://doi.org/10.1137/15M1033587)
718 [15M1033587](https://doi.org/10.1137/15M1033587).
- 719 Cabanac, G.; and Preuss, T. 2013. Capitalizing on order
720 effects in the bids of peer-reviewed conferences to secure
721 reviews by expert referees. *Journal of the Association for*
722 *Information Science and Technology* 64(2): 405–415.
- 723 Charlin, L.; and Zemel, R. S. 2013. The Toronto Paper
724 Matching System: An automated paper-reviewer assignment
725 system. In *ICML Workshop on Peer Reviewing and Publish-*
726 *ing Models*.
- 727 Chawla, S.; Dwork, C.; McSherry, F.; and Talwar, K. 2005.
728 On Privacy-Preserving Histograms. In *Conf. Uncertainty in*
729 *Artificial Intelligence (UAI)*, 120–127.
- Douceur, J. R. 2009. Paper rating vs. paper ranking. *ACM*
SIGOPS Operating Systems Review 43(2): 117–121.
- Dwork, C. 2006. Differential privacy. In *Proc. Int. Conf.*
Automata, Languages and Programming (ICALP), 1–12.
Venice, Italy.
- Dwork, C.; Kenthapadi, K.; McSherry, F.; Mironov, I.; and
Naor, M. 2006a. Our data, ourselves: privacy via distributed
noise generation. In *Proc. Annu. Int. Conf. Theory and*
Applications of Cryptographic Techniques (EUROCRYPT),
486–503. St. Petersburg, Russia.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2006b.
Calibrating noise to sensitivity in private data analysis. In
Proc. Conf. Theory of Cryptography (TCC), 265–284. New
York, NY.
- Dwork, C.; McSherry, F.; Nissim, K.; and Smith, A. 2016.
Calibrating noise to sensitivity in private data analysis. *Jour-*
nal of Privacy and Confidentiality 7(3): 17–51.
- Dwork, C.; and Roth, A. 2014. The Algorithmic Founda-
tions of Differential Privacy. *Found. Trends Theor. Comput.*
Sci. 9(3–4): 211–407.
- Fiez, T.; Shah, N.; and Ratliff, L. 2019. A SUPER* Algo-
rithm to Optimize Paper Bidding in Peer Review. In *ICML*
workshop on Real-world Sequential Decision Making: Re-
inforcement Learning And Beyond.
- Gaboardi, M.; Arias, E. J. G.; Hsu, J.; Roth, A.; and Wu,
Z. S. 2014. Dual Query: Practical Private Query Release
for High Dimensional Data. In *Int. Conf. Machine Learning*
(ICML). Beijing, China.
- Ge, H.; Welling, M.; and Ghahramani, Z. 2013. A Bayesian
model for calibrating conference review scores. URL {[http://](http://mlg.eng.cam.ac.uk/hong/nipsrevcal.pdf)
mlg.eng.cam.ac.uk/hong/nipsrevcal.pdf}.
- Goldsmith, J.; and Sloan, R. H. 2007. The AI conference pa-
per assignment problem. In *Proc. AAAI Workshop on Prefer-*
ence Handling for Artificial Intelligence, Vancouver, 53–57.
- Hay, M.; Rastogi, V.; Miklau, G.; and Suciu, D. 2010. Boost-
ing the accuracy of differentially private histograms through
consistency. *Proceedings of the VLDB Endowment* 3(1-2):
1021–1032.
- Jecmen, S.; Zhang, H.; Liu, R.; Shah, N. B.; Conitzer, V.;
and Fang, F. 2020. Mitigating Manipulation in Peer Review
via Randomized Reviewer Assignments. *arXiv*.
- Kerr, S.; Tolliver, J.; and Petree, D. 1977. Manuscript char-
acteristics which influence acceptance for management and
social science journals. *Academy of Management Journal*
20(1): 132–141.
- Kliwer, M. A.; DeLong, D. M.; Freed, K.; Jenkins, C. B.;
Paulson, E. K.; and Provenzale, J. M. 2004. Peer review
at the American Journal of Roentgenology: How reviewer
and manuscript characteristics affected editorial decisions
on 196 major papers. *American Journal of Roentgenology*
183(6): 1545–1550.
- Kobren, A.; Saha, B.; and McCallum, A. 2019. Pa-
per Matching with Local Fairness Constraints. In *ACM*
SIGKDD International Conference on Knowledge Discov-
ery & Data Mining.

- 785 Lawrence, N.; and Cortes, C. 2014. The NIPS Experiment. <http://inverseprobability.com/2014/12/16/the-nips-experiment>. [Online; accessed 11-June-2018].
- 786
- 787
- 788 Lee, C. J. 2015. Commensuration bias in peer review. *Philosophy of Science* 82(5): 1272–1283.
- 789
- 790 Li, C.; Hay, M.; Rastogi, V.; Miklau, G.; and McGregor, A. 2010. Optimizing Linear Counting Queries Under Differential Privacy. In *Symp. Principles Database Systems (PODS)*, 123–134. Indianapolis, IN.
- 791
- 792
- 793
- 794 Lindsey, D. 1988. Assessing precision in the manuscript review process: A little better than a dice roll. *Scientometrics* 14(1-2): 75–82.
- 795
- 796
- 797 Mavrogenis, A. F.; Quaile, A.; and Scarlat, M. M. 2020. The good, the bad and the rude peer-review.
- 798
- 799 McCook, A. 2006. Is peer review broken? Submissions are up, reviewers are overtaxed, and authors are lodging complaint after complaint about the process at top-tier journals. What’s wrong with peer review? *The scientist* 20(2): 26–35.
- 800
- 801
- 802
- 803 Noothigattu, R.; Shah, N. B.; and Procaccia, A. D. 2018. Loss Functions, Axioms, and Peer Review. *arXiv preprint arXiv:1808.09057*.
- 804
- 805
- 806 Peters, D. P.; and Ceci, S. J. 1982. Peer-review practices of psychological journals: The fate of published articles. *Behavioral and Brain Sciences* 5(2): 187–255.
- 807
- 808
- 809 Reinhart, M. 2009. Peer review of grant applications in biology and medicine. Reliability, fairness, and validity. *Scientometrics* 81(3): 789–809.
- 810
- 811
- 812 Rennie, D. 2016. Make peer review scientific. *Nature* 535(7610): 31–34.
- 813
- 814 Roos, M.; Rothe, J.; and Scheuermann, B. 2011. How to Calibrate the Scores of Biased Reviewers by Quadratic Programming. In *AAAI Conference on Artificial Intelligence*.
- 815
- 816
- 817 Schroter, S.; Loder, E.; and Godlee, F. 2020. Research on peer review and biomedical publication.
- 818
- 819 Scott, W. A. 1974. Interreferee agreement on some characteristics of manuscripts submitted to the Journal of Personality and Social Psychology. *American Psychologist* 29(9): 698.
- 820
- 821
- 822
- 823 Shah, N. B.; Tabibian, B.; Muandet, K.; Guyon, I.; and Von Luxburg, U. 2018. Design and analysis of the NIPS 2016 review process. *The Journal of Machine Learning Research* 19(1): 1913–1946.
- 824
- 825
- 826
- 827 Siegelman, S. S. 1991. Assassins and zealots: variations in peer review. *Radiology* 178(3): 637–642.
- 828
- 829 Snodgrass, R. 2006. Single- Versus Double-Blind Reviewing: An Analysis of the Literature. *SIGMOD Record* 35: 8–21. doi:10.1145/1168092.1168094.
- 830
- 831
- 832 Squazzoni, F.; Ahrweiler, P.; Barros, T.; Bianchi, F.; Birukou, A.; Blom, H. J.; Bravo, G.; Cowley, S.; Dignum, V.; Dondio, P.; et al. 2020. Unlock ways to share data on peer review.
- 833
- 834
- 835
- 836 Stelmakh, I.; Shah, N.; and Singh, A. 2018. PeerReview4All: Fair and Accurate Reviewer Assignment in Peer Review. *arXiv preprint arxiv:1806.06237*.
- 837
- 838
- Stelmakh, I.; Shah, N.; and Singh, A. 2019. On Testing for Biases in Peer Review. In *NeurIPS*. 839 840
- Stelmakh, I.; Shah, N.; and Singh, A. 2020. Catch Me if I Can: Detecting Strategic Behaviour in Peer Assessment. *arXiv*. 841 842 843
- Stelmakh, I.; Shah, N.; Singh, A.; and Daumé III, H. 2020. Prior and Prejudice: The Bias against Resubmissions in Conference Peer Review. *arXiv*. 844 845 846
- Tomkins, A.; Zhang, M.; and Heavlin, W. D. 2017. Reviewer bias in single-versus double-blind peer review. *Proceedings of the National Academy of Sciences* 114(48): 12708–12713. 847 848 849
- Wang, J.; and Shah, N. B. 2019. Your 2 is My 1, Your 3 is My 9: Handling Arbitrary Miscalibrations in Ratings. In *AAMAS*. 850 851 852
- Wang, W.; Ying, L.; and Zhang, J. 2016. On the relation between identifiability, differential privacy, and mutual-information privacy. *IEEE Trans. Inf. Theory* 62(9): 5018–5029. 853 854 855 856
- Welch, I. 2014. Referee recommendations. *The Review of Financial Studies* 27(9): 2773–2804. 857 858
- Xu, Y.; Zhao, H.; Shi, X.; and Shah, N. 2019. On Strategyproof Conference Review. In *IJCAI*. 859 860

861 **Appendices**

862 **A Algorithm for Upper Bounds**

863 The computation of upper bounds is similar to the method-
 864 ology in Algorithm 1 from Section 4.3 and is presented in
 865 Algorithm 2. The two criteria we use to determine if mean
 866 of a tuple is an upper bound on θ_i^* are
 867 C3 The longest right chain of the tuple has length at least
 868 $n - i + 1$.
 869 C4 In X , after we mark the ℓ weights from each tuple con-
 870 sidered so far, each row has at most $i - 1$ unmarked
 871 entries.
 872

Algorithm 2: Computation of upper bounds

Input: matrix X of weights, sorted list of weight tuples Ω
Initialize $i = 1$, set $w \in \mathbb{R}^\ell$ as the first tuple in Ω , and all
 entries of X are unmarked.
while $i \geq 1$ **do**
 for each weight in the tuple w , find its corresponding
 entry in matrix X and mark the entry
 if length of tuple w 's longest right chain $\geq n - i + 1$
 and number of unmarked entries on each row of
 $X \leq i - 1$ **then**
 upper bound on $\theta_i^* = \text{mean of all entries of tuple } w$
 $i- = 1$
 end if
 set w as the previous tuple in Ω
end while

873 **B An Example**

874 In this section, we illustrate our algorithm (described in Sec-
 875 tion 4.3) by means of an illustrative example.
 876 Consider a case where $n = m = 4$, $\ell = k = 3$. Let 3
 877 papers among the 4 have all 0 weights and the fourth paper
 878 has weights 1, 2 and 3. In this example, we can infer that
 879 $(\theta_1^*, \theta_2^*, \dots, \theta_n^*) = (0, \frac{1}{3}, \frac{2}{3}, 1)$ regardless of assignment. This
 880 example reflects axiomatic property A3 presented in Section
 881 4.3. We show that our algorithm for computing bounds
 882 indeed results in a convex set that contains only the vector
 883 $(0, \frac{1}{3}, \frac{2}{3}, 1)$. And thus the projection of any noisy data onto
 884 this convex set results in $(0, \frac{1}{3}, \frac{2}{3}, 1)$.
 885 First, we visualize the matrix X as

$$\begin{array}{l} \text{paper 1 : } 0_{11} \quad 0_{12} \quad 0_{13} \\ \text{paper 2 : } 0_{21} \quad 0_{22} \quad 0_{23} \\ \text{paper 3 : } 0_{31} \quad 0_{32} \quad 0_{33} \\ \text{paper 4 : } 1_{41} \quad 2_{42} \quad 3_{43}. \end{array}$$

886
 887 The subscripts indicate the entries of the weights
 888 in X . Some elements in Ω' are: $(0_{11}, 0_{21}, 0_{31})$,
 889 $\dots, (0_{13}, 0_{23}, 0_{33})$, $(0_{11}, 0_{12}, 1_{41})$, $\dots, (0_{13}, 0_{23}, 1_{41})$,
 890 $(0_{11}, 0_{12}, 2_{42})$, $\dots, (0_{11}, 0_{12}, 3_{43})$. After we sort Ω' based
 891 on the mean of the weight tuples, we get Ω where the first
 892 27 tuples have mean 0, followed by 27 tuples with mean
 893 $\frac{1}{3}$, 27 tuples with mean $\frac{2}{3}$, and 27 tuples with mean 1. We
 894 construct graph G in which for instance, there is an edge

between the tuples $(0_{11}, 0_{21}, 0_{31})$ and $(0_{12}, 0_{22}, 0_{32})$ since
 all six weights in these two tuples correspond to different
 entries in X . On the other hand, there is no edge between
 the tuples $(0_{11}, 0_{21}, 1_{41})$ and $(0_{12}, 0_{22}, 1_{41})$ because they
 both contain weight 1_{41} .

The first tuple in Ω meets both the criteria so lower bound
 on θ_1^* is mean of the first tuple, which is a tuple with mean
 0. Therefore, lower bound on θ_1^* is 0. Without loss of gener-
 ality, the first tuple is $(0_{11}, 0_{21}, 0_{31})$ and we mark the corre-
 sponding entries in X . Now the matrix X can be visualized
 as follows (where we mark any entry when we need to):

$$\begin{array}{l} \text{paper 1 : } \cancel{0}_{11} \quad 0_{12} \quad 0_{13} \\ \text{paper 2 : } \cancel{0}_{21} \quad 0_{22} \quad 0_{23} \\ \text{paper 3 : } \cancel{0}_{31} \quad 0_{32} \quad 0_{33} \\ \text{paper 4 : } 1_{41} \quad 2_{42} \quad 3_{43}. \end{array}$$

To compute a lower bound on θ_2^* , we start from the second
 tuple in Ω . Since there are 27 tuples that have mean zero, the
 second tuple still has mean zero. However, we do not choose
 any tuple that has mean zero due to criterion C2 from the
 algorithm. Choosing any $(0, 0, 0)$ tuple leaves all 3 entries
 of row 4 in X unmarked, and thus will not leave row 4 with
 at most 2 unmarked entries. Therefore, we will only stop at
 the first tuple that has mean $\frac{1}{3}$. Without loss of generality,
 we choose tuple $(0_{11}, 0_{21}, 1_{41})$ and mark the corresponding
 entries in X . This will leave the matrix X as

$$\begin{array}{l} \text{paper 1 : } \cancel{0}_{11} \quad \cancel{0}_{12} \quad \cancel{0}_{13} \\ \text{paper 2 : } \cancel{0}_{21} \quad \cancel{0}_{22} \quad \cancel{0}_{23} \\ \text{paper 3 : } \cancel{0}_{31} \quad \cancel{0}_{32} \quad \cancel{0}_{33} \\ \text{paper 4 : } \cancel{1}_{41} \quad 2_{42} \quad 3_{43}. \end{array}$$

For a similar reason, we do not choose any tuple that has
 mean $\frac{1}{3}$ to be a lower bound on θ_3^* as it would not leave row
 4 of X with at most 1 unmarked entry. So we choose the first
 tuple that has mean $\frac{2}{3}$ and a lower bound on θ_3^* is $\frac{2}{3}$. Lastly, a
 lower bound on θ_4^* is computed using the first tuple that has
 mean 1.

Now we can look at the computation of upper bounds us-
 ing the proposed algorithm. Upper bound on θ_4^* is taken as
 mean of the last tuple, which is a tuple with mean 1. There-
 fore, an upper bound on θ_4^* is 1. In addition, we mark two
 entries of weight 0 and one entry of weight 3. Without loss
 of generality, we mark entries $0_{11}, 0_{21}, 3_{43}$. Now the matrix
 X can be visualized as

$$\begin{array}{l} \text{paper 1 : } \cancel{0}_{11} \quad 0_{12} \quad 0_{13} \\ \text{paper 2 : } \cancel{0}_{21} \quad 0_{22} \quad 0_{23} \\ \text{paper 3 : } 0_{31} \quad 0_{32} \quad \cancel{0}_{33} \\ \text{paper 4 : } 1_{41} \quad 2_{42} \quad \cancel{3}_{43}. \end{array}$$

To compute an upper bound on θ_3^* , we start from the second
 to last tuple in Ω . Since there are 27 tuples that have mean
 1, the second to last tuple still has mean 1. However, we do
 not choose any tuple with mean 1 due to criterion C3 from
 the algorithm. Any $(0, 0, 3)$ tuple does not have a right chain
 of length at most 2 because all tuples with value $(0, 0, 3)$
 are not connected due to the uniqueness of the weight 3.
 Therefore, we will only stop at the first tuple that has mean $\frac{2}{3}$
 as it has a right chain of length 2. Since we have encountered
 all combinations of $(0, 0, 3)$, the matrix X after we choose a
 tuple $(0, 0, 2)$ becomes

paper 1 : $\frac{0}{11}$ $\frac{0}{12}$ $\frac{0}{13}$
 paper 2 : $\frac{0}{21}$ $\frac{0}{22}$ $\frac{0}{23}$
 paper 3 : $\frac{0}{31}$ $\frac{0}{32}$ $\frac{0}{33}$
 paper 4 : $\frac{1}{41}$ $\frac{2}{42}$ $\frac{3}{43}$

For a similar reason, we do not choose any tuple that has mean $\frac{2}{3}$ to be an upper bound on θ_2^* as it would not have a right chain of length at least 3. So we choose the first tuple we encounter that has mean $\frac{1}{3}$ and an upper bound on θ_2^* is $\frac{1}{3}$. Lastly, an upper bound on θ_1^* is computed using the first tuple that has mean 0.

Thus, the bounds on $\theta^* = (\theta_1^*, \theta_2^*, \dots, \theta_n^*)$ are $0 \leq \theta_1^* \leq 0$, $\frac{1}{3} \leq \theta_2^* \leq \frac{1}{3}$, $\frac{2}{3} \leq \theta_3^* \leq \frac{2}{3}$ and $1 \leq \theta_4^* \leq 1$. Along with the conditions that $\theta_1^* + \theta_2^* + \theta_3^* + \theta_4^* = 2$ and $\theta_1^* \leq \theta_2^* \leq \theta_3^* \leq \theta_4^*$, the only possible value of θ^* is $(0, \frac{1}{3}, \frac{2}{3}, 1)$. Thus the output of our algorithm is the singleton set $\{(0, \frac{1}{3}, \frac{2}{3}, 1)\}$. Projection of any data to the convex set $\{(0, \frac{1}{3}, \frac{2}{3}, 1)\}$ results in $(0, \frac{1}{3}, \frac{2}{3}, 1)$, which is consistent with axiomatic property A3.

C Simulations

In this section, we conduct synthetic simulations to evaluate the performance of our algorithm. We synthetically generate a conference review setting in one of several ways as described below. In each of the settings, the number of reviewers equals the number of papers, and each reviewer reviews 2 papers and each paper is reviewed by two reviewers. The assignment of reviewers to papers is done uniformly at random subject to given load constraints. The weight given by any reviewer to any reviewed paper is drawn from a beta distribution. For preserving privacy, we consider the common method of adding i.i.d. Laplace noise (with mean zero and variance 2) to each component of the true sorted mean-weight vector.

We evaluate the following three methods of releasing the sorted mean-weight vector, which includes our proposed algorithm and two baselines:

- **Noisy** where Laplace noise is added but no post-processing is performed;
- **Baseline projection** where the noisy data is post-processed via projecting onto a convex set which constrains the sum of all entries, the value of each entry in terms of the range of weights (0 to 1), and imposes a monotonicity constraint; We project on the set $\{\mathbf{t} \in \mathbb{R}^n | 0 \leq t_i \leq 1 \forall i \in [n], \sum_{i=1}^n t_i = \frac{1}{\ell} \sum_{i=1}^m \sum_{j=1}^k x_{ij}, t_1 \leq t_2 \leq \dots \leq t_n\}$.

• **Our algorithm** where the noisy data is post-processed via our algorithm described in Section 4.

The simulations compute the mean squared error between the true sorted mean-weight vector θ^* and the output from each of these three methods, that is, $\sum_{i=1}^n (t_i - \theta_i^*)^2$ where \mathbf{t} is the output of any of these algorithms. Note that in the figures, the error bars (standard error of the mean) are plotted but not visible in most cases since they are too small.

We now describe the method for generating the weights in each simulation, and refer the reader to the corresponding plots. Note that the y-axes (representing the mean squared error) on each of the plots is on a logarithmic scale.

- In Figure 2a—2e, the number of reviewers ranges from 10 to 50. The weights are all i.i.d. and are generated from the beta distribution specified in the corresponding sub-caption.
- In Figure 2f, the number of reviewers is fixed at 10. On the x-axis, we vary a parameter $a \in \{0.5, 1, \dots, 10\}$. For each value of a , we draw all weights i.i.d. from the $\text{beta}(a, a)$ distribution.
- In Figure 2g, we again vary the number of reviewers n on the x-axis. For any paper $i \in [n]$, one weight is generated from $\text{beta}(1, i)$ and the other weight is generated from $\text{beta}(2, i)$ independently.
- In Figure 2h, whenever any paper $i \in [m]$ is reviewed by reviewer $j \in [n]$, the weight of that review is generated from $\text{beta}(i, j)$.

All in all, these simulations reveal that our algorithm can lead to a multi-fold improvement in the utility (accuracy) while not compromising the privacy.

D Proofs

We present proofs of all the claimed results.

D.1 Proof of Proposition 4.1

We prove the proposition using a counter example. Assume the true value $\theta^* = 0$ and the set of all possible values $\Theta = \{-4, -2, 0, 2, 4\}$. The noisy data $\mathbf{r} = \theta^* + \eta$ where η is a Laplace random variable with probability density function $\eta(x) = 0.5e^{-|x|}$.

Without projection, the expected error incurred by the noise is $\int_{-\infty}^{\infty} 0.5e^{-|x|}x^2dx = 2$. But if we project the noisy data on the set Θ and get result \mathbf{t} , the expected error after the projection is computed as $16 \int_{-\infty}^{-3} 0.5e^{-|x|}dx + 4 \int_{-3}^{-1} 0.5e^{-|x|}dx + 4 \int_{1}^3 0.5e^{-|x|}dx + 16 \int_3^{\infty} 0.5e^{-|x|}dx = 2.06896$, which is greater than the expected error without projection. Thus, projecting on the set that contains all true values could decrease the accuracy of data.

D.2 Proof of Proposition 4.2

It is known that projection on a closed convex set is non-expansive (Bauschke, Combettes et al. 2011). Since θ^* results from a valid assignment, it is contained in Θ . Therefore it is contained in any closed convex set that contains Θ . Projection of \mathbf{r} onto any such convex set will not increase its squared error from θ^* . Therefore, proposition 4.2 holds.

D.3 Proof of Theorem 4.3

We will prove the NP-hardness by reducing the ℓ -Partition problem, which is NP-hard (Babel, Kellerer, and Kotov 1998), to the problem of projecting noisy data onto convex hull of Θ . The ℓ -Partition problem where $\ell > 2$ is defined as follows.

Definition D.1. ℓ -Partition problem: Given a multi-set $\mathcal{W} = \{w_1, w_2, \dots, w_n\}$ of n non-negative integers where n is a multiple of ℓ , decide if we can partition \mathcal{W} into $\frac{n}{\ell}$ subsets such that each subset has size ℓ and the sums of all subsets are the same.

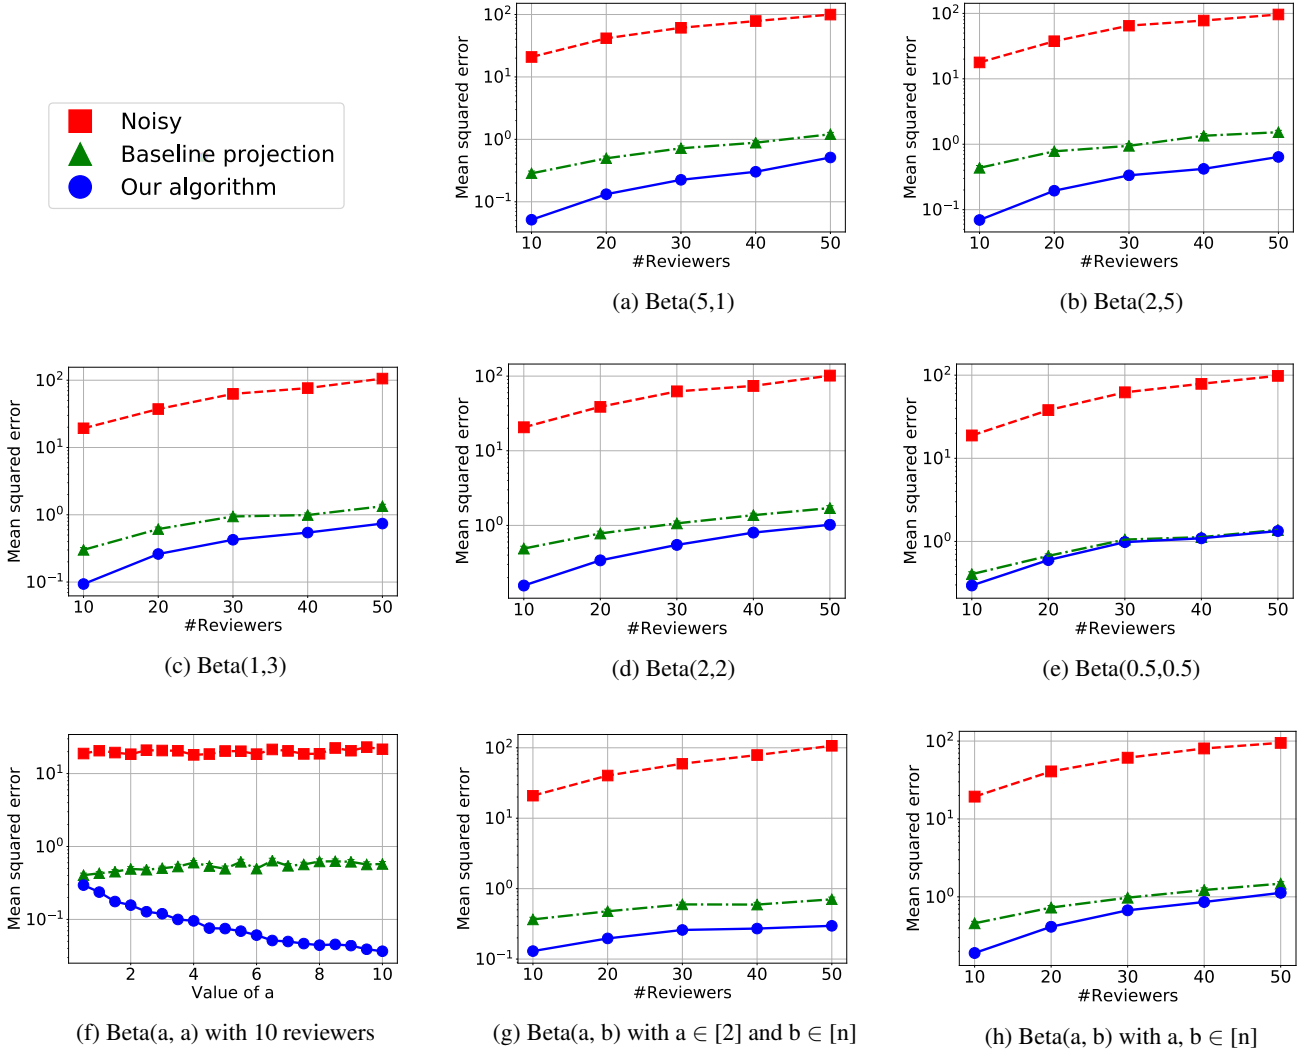


Figure 2: Simulation results. The y-axes of all plots are on a logarithmic scale.

1047 Consider any instance of the ℓ -Partition problem with
 1048 $\mathcal{W} = \{w_1, w_2, \dots, w_n\}$, where $w_i \geq 0$ and n is a multiple of
 1049 ℓ . Now we construct a peer-review dataset where there are
 1050 n reviewers and n papers, each reviewer reviews ℓ papers
 1051 and each paper receives ℓ reviews. Note that the number of
 1052 reviewers is the same as the number of elements in \mathcal{W} . Let
 1053 each paper i has weight w_i and $\ell - 1$ zero weights, and a
 1054 denote the average of all elements in \mathcal{W} , i.e.,

$$a = \frac{1}{n} \sum_{i=1}^n w_i. \quad (\text{D.1})$$

1055 Let $v = (0, \dots, 0, a, \dots, a)$ be a vector of n entries whose last
 1056 $\frac{n}{\ell}$ entries all have value a . Let $\mathcal{V} = \mathcal{W} \cup \{0, \dots, 0\}$ be the
 1057 multiset containing all values of \mathcal{W} and $n \cdot (\ell - 1)$ zeros.
 1058 Then the projection problem is to project v onto the convex
 1059 hull of Θ defined for this peer-review dataset.

1060 The reduction from the ℓ -Partition problem to the projec-
 1061 tion problem constructed above is as follows. If the solution

to the projection problem is v itself, we return True for the
 ℓ -Partition problem; otherwise we return False.

1062 We first prove the correctness of the reduction. Suppose
 1063 \mathcal{W} can be ℓ -partitioned into $\frac{n}{\ell}$ subsets of equal sums. Then
 1064 we can partition \mathcal{V} into subsets of size ℓ where these subsets
 1065 are the subsets that give the ℓ -partition of \mathcal{W} and subsets
 1066 that consist of ℓ zeros. By Lemma D.2 below, this partition
 1067 of \mathcal{V} gives a valid assignment for the peer-review problem,
 1068 and thus $v = (0, \dots, 0, a, \dots, a)$ corresponds to a valid as-
 1069 signment. Therefore, the projection of v is itself. The proof
 1070 of Lemma D.2 is presented in Section D.4.
 1071

Lemma D.2. *In the setting described above, any ℓ -partition
 of the $n \cdot \ell$ weights in \mathcal{V} , i.e., any partition of \mathcal{V} into subsets
 of size ℓ , can be interpreted as a valid assignment such that
 subset i corresponds to the weights from reviewer i given to
 ℓ distinct papers.*

1072 Next, suppose that the projection of v is itself. We show
 1073 that \mathcal{W} can be ℓ -partitioned into subsets of equal sums. We
 1074 first claim that v must correspond to a valid assignment it-
 1075
 1076
 1077
 1078
 1079
 1080

1081 self. To see this, suppose $v = (0, \dots, 0, a, \dots, a)$ is a con- 1139
 1082 vex combination of some sorted mean weight vectors. Then 1140
 1083 these vectors must all have value a for their last $\frac{n}{\ell}$ entries 1141
 1084 since each of these mean weight vector is sorted. Due to 1142
 1085 the sum constraint, these mean weight vectors have to be 1143
 1086 $(0, \dots, 0, a, \dots, a)$. Next we note that in the assignment given 1144
 1087 by v , each reviewer who has an average weight of a must 1145
 1088 give ℓ weights with values from \mathcal{W} due to the pigeonhole 1146
 1089 principle. Therefore, the $\frac{n}{\ell}$ subsets each of which consists 1147
 1090 of weights given by one of the last $\frac{n}{\ell}$ reviewers form an ℓ - 1148
 1091 partition of \mathcal{W} with equal sums. 1149

1092 Finally, we prove the efficiency of the reduction. Since 1150
 1093 the construction of v has $\mathcal{O}(n)$ time complexity and the con- 1151
 1094 struction of \mathcal{V} has $\mathcal{O}(1)$ time complexity, the reduction has 1152
 1095 $\mathcal{O}(n)$ time complexity, which is polynomial in the size of 1153
 1096 the input. Thus, the reduction can be done efficiently, which 1154
 1097 completes the proof. 1155

1098 D.4 Proof of Lemma D.2

1099 Fix ℓ , we will prove the lemma by induction on n , the num- 1158
 1100 ber of reviewers, which is the same as the number of papers. 1159

1101 Base case: when $n = \ell$, every reviewer reviews all papers, 1160
 1102 so any ℓ partition of the weights can be validly assigned to 1161
 1103 reviewers. 1162

1104 Inductive hypothesis: suppose when there are fewer than 1163
 1105 n reviewers for an $n > \ell$, every ℓ partition of the weights 1164
 1106 forms a valid assignment for reviewers. 1165

1107 Consider when there are n reviewers and n papers. With- 1166
 1108 out loss of generality, assume all weights in \mathcal{W} are non-zero. 1167
 1109 Consider an ℓ partition of the set \mathcal{V} . We will argue in two 1168
 1110 cases based on whether there is a subset that contains ex- 1169
 1111 actly one non-zero weight. 1170

1112 1. Case 1: In the partition, if there is a subset with exactly 1171
 1113 one non-zero weight. 1172

1114 Without loss of generality, assume that the subset with 1173
 1115 exactly one non-zero weight contains w_1 and the subset 1174
 1116 is $\{w_1, 0, \dots, 0\}$. We denote the subset \mathcal{S}_1 . In \mathcal{S}_1 , there are 1175
 1117 $\ell - 1$ zero weights and a non-zero weight w_1 from \mathcal{W} . 1176
 1118 Since paper 1 receives ℓ weights in total, we can remove 1177
 1119 \mathcal{S}_1 , paper 1 and reviewer 1. 1178

1120 Now we are left with $n - 1$ reviewers and papers. The re- 1179
 1121 moval does not affect the number of reviews received by 1180
 1122 the rest of the papers. We still have each paper getting ℓ 1181
 1123 weights. Among the weights, there is one non-zero weight 1182
 1124 from \mathcal{W} and $\ell - 1$ zero weights. By the inductive hypoth- 1183
 1125 esis, the rest of the subsets in the partition form a valid 1184
 1126 assignment of $\mathcal{V} \setminus \mathcal{S}_1$. We can assign the weights to $n - 1$ 1185
 1127 reviewers. 1186

1128 We then add \mathcal{S}_1 back to the assignment. Since reviewer 1 1187
 1129 ℓ weights to paper 1, it is not valid. We can solve this by 1188
 1130 swapping the zero weights in \mathcal{S}_1 with zeros in other sub- 1189
 1131 sets. We need to make $\ell - 1$ swaps. We label the rest of the 1190
 1132 subsets $\mathcal{S}_2, \dots, \mathcal{S}_n$ where \mathcal{S}_2 is the subset that contains 1191
 1133 most non-zero weights and the labels go in decreasing or- 1192
 1134 der based on the number of non-zero weights contained in 1193
 1135 a subset. We look at the rest of the subsets in the order of 1194
 1136 their labels. 1195

1137 Since none of the rest of the subsets contain any weight 1196
 1138 from paper 1, swapping a zero weight from paper 1 into 1197

1139 any of these subsets will nor affect the validity of the sub- 1140
 1141 set. There are at least $n - 1 - \frac{n-1}{\ell}$ subsets that con- 1142
 1143 tain at least a zero weight. Since $n > \ell$ and $\ell > 2$, 1144
 $n - 1 - \frac{n-1}{\ell} = \frac{(\ell-1)(n-1)}{\ell} \geq \ell - 1$. Thus, we have 1145
 1146 enough subsets to swap the zero weights from paper 1 in. 1147
 1148 Then we make sure the zero weights swapped into \mathcal{S}_1 will 1149
 1149 not come from the same paper. We label the zeros in \mathcal{S}_1 1150
 1151 with index $1, \dots, \ell - 1$. Suppose there are no subsets that 1151
 1152 do not contain any zero weights. Then when we need to 1152
 1153 swap out the zero weight at index i in \mathcal{S}_1 , there are at most 1153
 1154 $n - i$ non-zero weights in the untouched subsets due to the 1154
 1155 order we look at the subsets. There are $n - i$ untouched 1155
 1156 subsets as well. Then there exists an untouched subset that 1156
 1157 contains i zero weights. Since at this stage \mathcal{S}_1 has already 1157
 1158 completed $i - 1$ swaps, we can find a zero weight from 1158
 1159 the untouched subset to swap so that the zero weight does 1159
 1160 not come from the same paper as the zero weights from 1160
 1161 previous swaps. Note that if we have any subset that does 1161
 1162 not contain any zero weight or we skip some subsets due 1162
 1163 to conflict of papers, then the fraction of non-zero weights 1163
 1164 left and untouched subsets will be even smaller. So we are 1164
 1165 guaranteed to find a proper zero weight to swap. Thus, we 1165
 1166 can make $\ell - 1$ swaps of the zero weights to \mathcal{S}_1 and make 1166
 1167 all subsets valid assignments of weights. Such swaps do 1167
 1168 not affect the values in each subset. 1168

1169 Therefore, such partition can result in a valid assignment 1169
 1170 of the $n \cdot \ell$ scores among n reviewers. 1170

2. Case 2: In the partition, if there are no subsets with exactly 1171
 1172 one non-zero weight. 1172

1173 Since \mathcal{W} contains n elements and there are n subsets, by 1173
 1174 pigeon hole principle, there must be a subset \mathcal{S}_1 that con- 1174
 1175 tains all zero weights. 1175

1176 Without loss of generality, we find the subset that con- 1176
 1177 tains w_1 and then swap w_1 with a zero weight in \mathcal{S}_1 . This 1177
 1178 results in $\mathcal{S}'_1 = \{w_1, 0, \dots, 0\}$. 1178

1179 Now we have a subset that contains exactly 1 weight from 1179
 1180 \mathcal{W} . Like in case 1, we remove the subset, reviewer 1 and 1180
 1181 paper 1. We can find a valid assignment of the rest of the 1181
 1182 weights to $n - 1$ reviewers. Then we will put \mathcal{S}'_1 back to 1182
 1183 the assignment. Currently all weights in \mathcal{S}'_1 are from paper 1183
 1184 1. We identify the subset where w_1 comes from, and swap 1184
 1185 w_1 back into the subset with a zero weight there. Since 1185
 1186 the subset can not contain any weights from paper 1, we 1186
 1187 can safely put w_1 back without having two weights from 1187
 1188 the same paper. 1188

1189 After the swap, \mathcal{S}'_1 has $\ell - 1$ zero weights from paper 1189
 1190 1 and a zero weight from a different paper, say paper 2. 1190
 1191 We need to make $\ell - 2$ swaps for the zeros in \mathcal{S}'_1 . We 1191
 1192 label the rest of the subsets $\mathcal{S}_2, \dots, \mathcal{S}_n$ where \mathcal{S}_2 is the 1192
 1193 subset that contains most non-zero weights and the labels 1193
 1194 go in decreasing order based on the number of non-zero 1194
 1195 weights contained in a subset. We look at the rest of the 1195
 1196 subsets in the order of their labels. 1196

1197 Since none of the rest of the subsets contain any weight 1197
 1198 from paper 1, swapping a zero weight from paper 1 into 1198
 1199 any of these subsets will nor affect the validity of the sub- 1199
 1200 set. In the worst case, there exists a subset that contains 1200
 1201 w_1 and there are at most $\frac{n-2}{\ell-1}$ subsets that only contains a 1201

1197 zero weight from paper 2 because such tuples cannot contain
 1198 w_2 . Then there are at least $n - 1 - \frac{n-2}{\ell-1} - 1$ subsets
 1199 that we can swap the zero weights in \mathcal{S}'_1 . Since $n > \ell$ and
 1200 $\ell > 2$, $n - 1 - \frac{n-2}{\ell-1} - 1 = \frac{(\ell-2)(n-2)}{\ell-1} \geq \ell - 2$. Thus, we
 1201 have enough subsets to swap the zero weights from paper
 1202 1 in.

1203 We keep a zero weight from paper 1 in \mathcal{S}_1 and label the
 1204 rest of the zero weights in \mathcal{S}_1 with index $1, \dots, \ell - 2$.
 1205 Suppose there are no subsets that do not contain any zero
 1206 weights. Then when we need to swap the zero weight at
 1207 index i in \mathcal{S}_1 , there are at most $n - i$ non-zero weights
 1208 in the untouched subsets due to the order we look at the
 1209 subsets. There are $n - i$ untouched subsets as well. Then
 1210 there exists a subset that contains $i + 1$ zero weights. Since
 1211 at this stage \mathcal{S}_1 has already completed $i - 1$ swaps, we
 1212 can find a zero weight to swap that does not conflict with
 1213 the weights from previous swaps and not from paper 2 either.
 1214 Note that if we have any subset that does not contain
 1215 any zero weight or we skip some subsets due to conflict
 1216 of papers, then the fraction of non-zero weights left and
 1217 untouched subsets will be even smaller. So we are guar-
 1218 anteed to find a proper zero weight to swap. Thus, we can
 1219 make $\ell - 2$ swaps of the zero weights to \mathcal{S}'_1 and makes all
 1220 subsets valid assignments of weights. Such swaps do not
 1221 affect the value in each subset.

1222 Therefore, such partition can result in a valid assignment
 1223 of the $n \cdot \ell$ weights among n reviewers.

1224 In conclusion, any ℓ -partition of \mathcal{V} can be interpreted as a
 1225 valid assignments of weights to n reviewers.

1226 D.5 Proof of Theorem 4.6

1227 We would like to show that the convex set contains Θ . We
 1228 will show that the bounds are indeed lower and upper bounds
 1229 on each entry.

1230 We will first show that the lower bounds computed by the
 1231 algorithm are correct.

1232 Assume for the sake of contradiction, there exists an as-
 1233 signment such that θ_i^* is less than the lower bound on θ_i^* we
 1234 computed, denoted as θ_i . We use ν to denote the tuple that
 1235 results in θ_i^* and use ν' to denote the tuple that we choose
 1236 in the algorithm that has mean θ_i . Since ν is a valid assign-
 1237 ment, it is the sum of ℓ weights from ℓ distinct papers. Since
 1238 Ω contains all such tuples, it contains ν . And since $\theta_i^* < \theta_i$,
 1239 we encountered ν before we encounter ν' in Ω . We did not
 1240 choose ν as the tuple for lower bound due to its violation of
 1241 either criterion C1 or criterion C2.

1242 If ν violates criterion C1, it does not have a left chain
 1243 of size at least i . There cannot be $i - 1$ weight tuples each
 1244 containing ℓ weights from different papers such that they
 1245 all have mean no larger than θ_i^* . Otherwise they form a left
 1246 chain of length i . So ν cannot have its mean appear at entry
 1247 i in θ^* .

1248 If ν violates criterion C2, there exists a row that has more
 1249 than $n - i$ unmarked entries in X . The weights of the un-
 1250 marked entries have not been encountered so far, which in-
 1251 dicates that any tuple that contains the weights from un-
 1252 marked entries has mean no less than θ_i^* . Otherwise, we
 1253 would have encountered the weight before ν and mark its

entry. We know that there are $n - i$ reviewers who has mean
 weight no less than θ_i^* . In addition, there are more than $n - i$
 weights left for at least one paper. By Pigeon Hole Principle,
 there exists a reviewer gives a weight tuple that contains two
 weights from the same paper. However, no two weights from
 the same paper can be in the same tuple since one reviewer
 cannot give 2 weights to the same paper. So ν cannot have
 its mean appear at entry i in θ^* .

Thus, θ_i^* cannot be a value for entry i in θ^* . The value θ_i
 we computed is indeed a lower bound on that entry.

Following a similar argument, we can prove the correct-
 ness of the upper bounds from the algorithm.

1266 D.6 Proof of Theorem 4.7

1267 We will show that the proposed algorithm has polynomial
 1268 time complexity in the number of reviewers. There are
 1269 $n \cdot \ell$ weights, so the size of Ω' , denoted $|\Omega'|$, has size at
 1270 most $\binom{n \cdot \ell}{\ell}$, which is of complexity $\mathcal{O}(n^\ell)$. Sorting Ω' has
 1271 $\mathcal{O}(|\Omega'| \log(|\Omega'|))$ time complexity, which is still polynomial
 1272 in n . There are $\binom{|\Omega'|}{2}$ pairs of vertices to examine for edges.
 1273 Therefore, constructing G is of polynomial time in n . To
 1274 compute the length longest left chain and right chain of a
 1275 vertex, we can make use of a dynamic programming algo-
 1276 rithm that only requires us to loop through Ω once to com-
 1277 pute length of longest left chain of all vertices and loop one
 1278 more time to compute the length of longest right chain. For
 1279 each vertex, we examine at most all its neighbors, which is
 1280 of size polynomial in n . Lastly, after all preparation work,
 1281 for each vertex, we take $\mathcal{O}(1)$ time to check criteria C1 and
 1282 and C3 at most $\mathcal{O}(m)$ time to check criteria C2 and C4.
 1283 Since $m \leq n \cdot \ell$, both operations are polynomial in n . Thus,
 1284 the proposed algorithm computes the bounds in time poly-
 1285 nomial in n .

We will use quadratic programming to project noisy data
 onto the convex set and there are $2n$ linear constraints. This
 operation is also polynomial in n .

Thus, the proposed algorithm has time complexity that is
 polynomial in n .

1291 D.7 Proof of Theorem 4.8

1292 Axiomatic property A1: When all weights are the same, all
 1293 weight tuples have the same mean, which equals the weight.
 1294 Thus, all lower and upper bounds have the same value as the
 1295 weight. The convex set contains a single vector and projec-
 1296 tion of any noisy data on such convex set will result in the
 1297 vector, whose entries are all the same as the weight.

1298 Axiomatic property A2: When $\ell = 1$, there are exactly n
 1299 weight tuples, each containing one weight. We will choose
 1300 the same weight tuple for lower bound and upper bound on
 1301 θ_i^* . The mean of the chosen weight tuple is the weight of
 1302 rank i among all n weights. Therefore, the convex set con-
 1303 tains exactly one vector, which is the sorted vector of all
 1304 weights. Projection of any noisy data onto this convex set
 1305 will result in the vector of sorted weights.

1306 Axiomatic property A3: When all except for one paper re-
 1307 ceives all zero weights, computation of lower bound on θ_i^*
 1308 when $i < n - k$ will choose a tuple whose weights are all
 1309 zeros. When $i \geq n - k$, computation of lower bound will

1310 choose a tuple that contains a nonzero weights due to crite-
1311 rion C2. Similarly, to compute an upper bound on θ_i^* when
1312 $i \geq n - k$, we will choose a tuple with a nonzero weight
1313 due to the criterion C3. But when $i < n - k$, the algorithm
1314 will choose a tuple with all zero weights. The example we
1315 present in Section B illustrates this process. Therefore, the
1316 convex set again contains only a vector who has $n - k$ zero
1317 entries. Projection of any noisy data will result in this vector.