

Mitigating the Participation Bias by Balancing Extreme Ratings

Anonymous Author(s)*

Abstract

Rating aggregation plays a crucial role in various fields, such as product recommendations, hotel rankings, and teaching evaluations. However, traditional averaging methods can be affected by participation bias, where some raters do not participate in the rating process, leading to potential distortions. In this paper, we consider a robust rating aggregation task under the participation bias. We assume that raters may not reveal their ratings with a certain probability depending on their individual ratings, resulting in partially observed samples. Our goal is to minimize the expected squared loss between the aggregated ratings and the average of all underlying ratings (possibly unobserved) in the worst-case scenario.

We focus on two settings based on whether the sample size (i.e. the number of raters) is known. In the first setting, where the sample size is known, we propose an aggregator, named as the Balanced Extremes Aggregator. It estimates unrevealed ratings with a balanced combination of extreme ratings. When the sample size is unknown, we derive another aggregator, the Polarizing-Averaging Aggregator, which becomes optimal as the sample size grows to infinity. Numerical results demonstrate the superiority of our proposed aggregators to participation bias, compared to simple averaging.

CCS Concepts

• Theory of computation → Algorithmic game theory; Algorithmic mechanism design.

Keywords

Rating Aggregation, Participation Bias, Robust Aggregation

ACM Reference Format:

Anonymous Author(s). 2018. Mitigating the Participation Bias by Balancing Extreme Ratings. In *Proceedings of Make sure to enter the correct conference title from your rights confirmation email (Conference acronym 'XX)*. ACM, New York, NY, USA, 14 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Suppose you come back from a hotel stay that left you frustrated. The service was terrible, and the noise kept you up all night. The experience was bad enough that, as soon as you got home, you felt compelled to go online and leave a negative review. This is not something you normally do—in fact, the last time you stayed at a hotel that was just okay, you did not even bother to rate it.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference acronym 'XX, June 03–05, 2018, Woodstock, NY

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM ISBN 978-1-4503-XXXX-X/18/06 <https://doi.org/XXXXXXX.XXXXXXX>




Figure 1: The observed ratings of two hotels.




Figure 2: The true ratings of two hotels.

Now, imagine you're on the other side, trying to decide between two hotels on an online platform. Both have identical average ratings, but as shown in Figure 1, their rating distributions are quite different. Which one would you pick?

While many people rely on average ratings to make their choices, these scores can be misleading due to biases in the way they are collected. One such bias is participation bias, where the ratings do not represent the views of all guests—e.g., only the ones who felt strongly enough to leave feedback [5, 57].

At Hotel A, reviews are polarized: some guests, possibly those staying in the premium rooms, report excellent stays with 5-star ratings, while others leave very negative feedback. In contrast, Hotel B's reviews are consistently around 4 stars, indicating generally positive but not outstanding experiences. This suggests that Hotel A's displayed rating might be inflated, as guests with neutral opinions may have skipped reviewing, whereas Hotel B's balanced ratings offer a more accurate reflection. Figure 2 shows a possible way of the true distribution.

Such participation bias affects other areas too, like course evaluations, movie reviews, and product ratings on e-commerce platforms, where only extreme opinions tend to be represented [5, 25, 27, 57].

If we could estimate the likelihood of participation across different rater groups, we could apply traditional methods like importance sampling to approximate the true average rater ratings [45, 46, 49]. However, without data from non-participants, quantifying the precise extent of this bias remains challenging. Therefore, we need to use alternative metrics to evaluate the aggregation methods without full knowledge of the participation bias. Following the

field of robustness analysis [3], we employ a robust aggregation paradigm aimed at minimizing the worst-case error to mitigate the participation bias.

1.1 Problem Statement

For all positive integer k , $[k]$ denotes the set $\{1, 2, \dots, k\}$.

We aim to find a robust function f to aggregate the discrete ratings from n raters. Each rating $x_i \in [m]$ follows an underlying distribution $\mathbf{p} = (p_1, p_2, \dots, p_m)$ and we want to estimate the expectation $\mu = \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$. We assume the ratings are independent. Otherwise, the actual ratings may contain systematic biases, making it impossible for any aggregator to recover the true expectation even given all ratings. Under this assumption, the best estimator for μ is the empirical average of all ratings $\frac{1}{n} \sum_i x_i$.

However, due to the participation bias, raters will report their ratings with probabilities that depend on the rating values. Let $\mathbf{g} = (g_1, g_2, \dots, g_m)$ represent these probabilities, where g_r is the probability of a rater reporting her rating when it is r . We assume raters are homogeneous and have the same participation bias. As a result, only a subset of elements in \mathbf{X} is observed, which we denote by $\hat{\mathbf{X}}$. For each rater i :

$$\hat{x}_i = \begin{cases} x_i & \text{with probability } g_{x_i} \\ 0 & \text{otherwise} \end{cases}$$

where 0 indicates that the rating is unobserved. Let $D_{\mathbf{p}}$ denote the distribution generating \mathbf{X} and $D_{\mathbf{p}, \mathbf{g}}$ denote the distribution generating $\hat{\mathbf{X}}$. We also call $D_{\mathbf{p}, \mathbf{g}}$ the information structure.

Without any information on the likelihood of reporting, the worst-case scenario would be no one reports, resulting in a lack of any useful aggregator. Thus, we assume there exists a known parameter q , indicating the lowest reporting probability among raters. That is, $g_r \in [q, 1]$ for any rating r . When $q = 1$, there is no participation bias and as q decreases, the bias may increase accordingly. We use an example to illustrate our model.

EXAMPLE 1.1 (WHY DIRECT AVERAGE IS NOT GOOD). Consider a binary rating system where $x_i \in \{1, 2\}$ with equal probability. Here, let $g_1 = 1$ and $g_2 = q$, meaning that negative reviews ($x_i = 1$) are always observed, while positive reviews ($x_i = 2$) are observed only with probability q .

Under this setup, even with a large sample size n , the direct average of observed ratings will tend toward $\frac{1+2q}{1+q}$, which significantly underestimates the true expectation of $3/2$ when q is small.

We measure the performance of an aggregator by its expected squared error, defined as $L(f, \mathbf{p}, \mathbf{g}) = \mathbb{E}_{\hat{\mathbf{X}} \sim D_{\mathbf{p}, \mathbf{g}}} [(f(\hat{\mathbf{X}}) - \mu)^2]$. One objective could be finding f that minimizes the worst-case error: $\max_{\mathbf{p}, \mathbf{g}} L(f, \mathbf{p}, \mathbf{g})$. However, when the sample size n is small, selecting \mathbf{p} with high variances can lead to large errors even for the ideal aggregator who observes \mathbf{x} . Since we focus on addressing the errors attributable to uncertainties in \mathbf{g} , we adopt a robust approach and instead minimize the regret relative to the ideal aggregator who observes the full data \mathbf{x} , and outputs $\frac{1}{n} \sum_i x_i$:

$$\min_f \max_{\mathbf{p}, \mathbf{g}} \mathbb{E}_{\mathbf{X} \sim D_{\mathbf{p}}, \hat{\mathbf{X}} \sim D_{\mathbf{p}, \mathbf{g}}} [(f(\hat{\mathbf{X}}) - \mu)^2] - \left(\frac{1}{n} \sum_i x_i - \mu \right)^2.$$

For short, we denote $R(f, \mathbf{p}, \mathbf{g}) = \mathbb{E}_{\mathbf{X} \sim D_{\mathbf{p}}, \hat{\mathbf{X}} \sim D_{\mathbf{p}, \mathbf{g}}} [(f(\hat{\mathbf{X}}) - \mu)^2] - \left(\frac{1}{n} \sum_i x_i - \mu \right)^2$.

Because we have assumed the raters are homogeneous, observing $\hat{\mathbf{x}}$ is equivalent as observing the histograms $n_u, n_1, n_2, \dots, n_m$ where n_r counts the number of ratings of r for all $r \in \{1, \dots, m\}$, and n_u indicates the number of unobserved ratings. That is, the aggregator's input is $n_u, n_1, n_2, \dots, n_m$. This is applicable in the scenarios when the sample size n is known. For example, in the teaching evaluation, the number of students is known to the instructors but not all students will give feedback.

A variant with unknown n . However, in practice, we may not know n . For example, in the movie rating platforms, we cannot obtain the number of people watching the movies. Thus, we consider a variant of the above problem where n is unknown. In this case, the aggregation problem remains the same except that the aggregator's input is n_1, n_2, \dots, n_m .

1.2 Summary of Results

We explore the optimal aggregator in two settings, depending on the knowledge of the sample size n (i.e., the number of raters).

The sample size n is known. In Section 3, we assume n is known. We construct a lower bound for the regret by considering a mixture of two specific information structures. Based on the lower bound, we provide a new aggregator, the Balanced Extreme Aggregator (BEA), which is the best response to those two specific information structures. Intuitively, BEA estimates the expected ratings of those unreported raters based on the difference between n_1 and n_m , the counts of extreme ratings. Then BEA adjusts the observed rating average with the estimated unobserved rating average. Figure 3 illustrates the whole process. We illustrate its near-optimal performance numerically in various cases.

The sample size n is unknown. In Section 4, we assume n is unknown. Since we do not know how many unobserved ratings there are, the above aggregator is not applicable. Instead, we obtain a new aggregator, the Polarizing-Averaging Aggregator (PAA), and show its optimality when n goes to infinite. Furthermore, we numerically show that PAA performs well for a finite sample size. Figure 4 illustrates the whole process. We create two modified histograms from the original observed data. For the first histogram, we identify a threshold k_1 , and only keep q fraction of the counts for ratings above k_1 . In the second histogram, we identify a threshold k_2 , and only keep q fraction of the counts for ratings below k_2 . We then calculate the empirical mean for each of these adjusted histograms and output their average.

Section 5 shows the virtualization of our aggregators, which helps to establish some insights about the aggregators.

2 Related Literature

Rating Aggregation. Rating aggregation can be defined as the problem of integrating different ratings into a single representative value. Previous works aim to resist individual disturbance from random or malicious spammers [4, 6] by building a reputation system [14, 26, 40]. The basic idea of reputation system is weighted average where a rater's weight is measured by his reputation score




Figure 3: Illustration of the Balanced Extreme Aggregator (BEA) when the sample size n is known. BEA estimates the expected ratings of those unreported raters based on the difference between n_1 and n_m , the counts of extreme ratings. Then BEA adjusts the observed rating average with the estimated unobserved rating average.

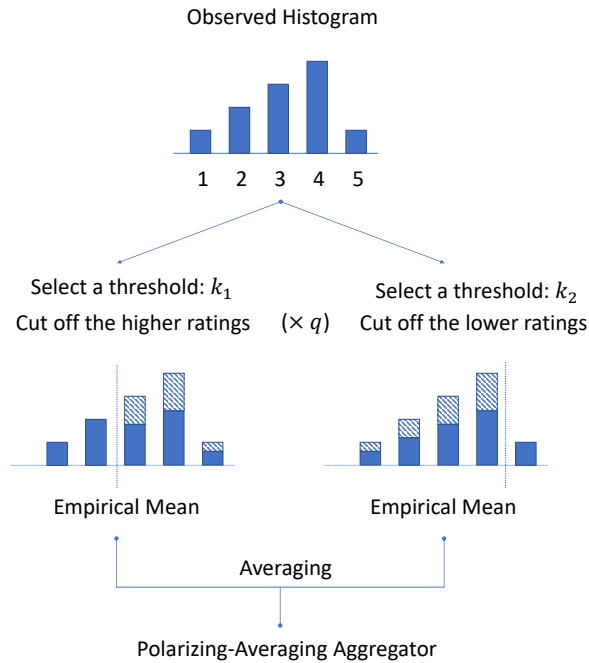


Figure 4: Illustration of the Polarizing-Averaging Aggregator (PAA) when the sample size n is unknown. We create two modified histograms from the original observed data. For the first histogram, we identify a threshold k_1 , and only keep q fraction of the counts for ratings above k_1 . In the second histogram, we identify a threshold k_2 , and only keep q fraction of the counts for ratings below k_2 . We then calculate the empirical mean for each of these adjusted histograms and output their average.

which is calculated by using the rater-object weighed bipartite network. The quality-based methods measure a raters's reputation by

the difference between the rating values and the objects' weighted average rating values [29, 33, 54, 56]. Unlike quality-based methods, raters are grouped based on their rating similarities and their reputation is calculated by the corresponding group sizes in group-based methods. [13, 15, 16] Another branch of literature considers how to resist collusive disturbance where malicious raters simultaneously promote or demote the qualities of the targeted objects by giving consistent rating score [36, 50–52, 55].

Our work focuses on the simple setting where there is a single object so we can not use the rater-object weighed bipartite network to learn raters' rating behavioral patterns. We only consider participation bias rather than other malicious behaviors. Finally, we adopt an adversarial approach to evaluate the performance of an aggregation scheme.

Participation Bias. Participation bias is a well-studied topic within the casual inference community which arises when participants disproportionately possess certain traits that affect participation [11]. It is a common source of error in clinical trials and survey studies [7, 8, 12, 42, 43].

When participation bias is only related to some observed (measured) variables and other representative samples are available (e.g. the whole population data from government), weighting strategies make it possible to create a pseudo-sample representative for these variables [41, 42] under the assumption that non-participants have equivalent behaviours to participants in the same socio-demographic category. An alternative way is multiple imputation (MI), which is viable when data are missing at random [2, 39, 44] and can be strengthened if we can infer information on the absent non-participants using record-linkage [17, 19] or other methods. If the data are thought to be missing-not-at-random then methods involving sensitivity analyses like pattern mixture modelling [34] and NARFCS [47] can be helpful. These methods can also be combined to correct participation bias [18].

Methods used to correct sample selection bias can also be used to correct participation bias [48].

The first solution to sample selection bias was suggested by Heckman [21] who proposed a maximum likelihood estimator which heavily relies on the normality assumption. One way to relax normality while remaining in the maximum likelihood framework was suggested by Lee [30, 31]. Two-Step estimators are frequently employed in empirical work including parametric models [22, 23] and semi-parametric models [1, 24, 32].

Recent works on sample selection models have aimed to address robust alternatives to the Heckman model. Marchenko and Genton [35] introduced a selection-t model, which models the error distribution using a Student's t distribution for dealing with the robustness to the normal assumption in the Heckman model. Zhelonkin et al. [53] proposed a modified robust semi-parametric alternative based on Heckman's two-step estimation method with asymptotic normality and gave its asymptotic variance. de Souza Bastos et al. [10] propose a generalization of the Heckman sample selection model by allowing the sample selection bias and dispersion parameters to depend on covariates.

Our work significantly differs from the above literature in a few aspects. First, we are only interested in recovering the average score rather than the casual effect. Second, we consider a simplified

setting where we only know the reported score and the lower bound of participation rate instead of various socio-demographic data. Finally, instead of making extra assumptions, we adopt a robust approach to evaluate the performance of an aggregation scheme.

Robust Aggregation. Arieli et al. [3] first introduced a robust paradigm for forecast aggregation, aiming to minimize the aggregator's regret compared to an omniscient aggregator. Later, De Oliveira et al. [9] explored a robust absolute paradigm focused on minimizing the aggregator's own loss. Since then, a growing number of researches have addressed robust information aggregation under various information structures, including the projective substitutes condition [37], conditional independent setting [3], and scenarios involving second-order information [38]. Furthermore, Guo et al. [20] provide an algorithmic approach for computing a near-optimal aggregator across general information structures. Kong et al. [28] tackle robust aggregation in the context of base rate neglect, a specific bias in forecast aggregation.

We consider a different setting from the above literature. We focus on addressing participation bias in rating aggregation.

3 Rating Aggregation with Known Sample Size

This section will discuss the setting with a known number of agents n . We provide the Balanced Extremes Aggregator (BEA), which performs near-optimal numerically in various cases.

BEA estimates the expectation of the unobserved ratings, $\hat{\mu}_u$, using a convex combination of extreme ratings, given by: $\hat{\mu}_u = \alpha \times 1 + (1 - \alpha) \times m$ where α is a parameter that depends on the observed ratings. Intuitively, α increases with the number of observed ratings of 1. BEA then combines $\hat{\mu}_u$ with the average observed ratings to produce a final aggregated result.

DEFINITION 3.1 (BALANCED EXTREMES AGGREGATOR). *Given sample size n , the observed ratings \hat{x} , the aggregator's output is defined as $f^{BEA}(\hat{x}) = \frac{n_o \hat{\mu}_o + n_u \hat{\mu}_u}{n}$. Here $n_o = \sum_i \mathbb{1}(\hat{x}_i \neq 0)$ is the number of observed ratings, $n_u = n - n_o$ is the number of unobserved ratings, $\hat{\mu}_o = \frac{\sum_i \hat{x}_i}{n_o}$ is the average of the observed ratings, $\hat{\mu}_u = \alpha \times 1 + (1 - \alpha) \times m$ is the estimation of the unobserved ratings' expectation. α^1 depends on the difference between the count of ratings of 1 and the count of ratings of m , denoted as $n_1 - n_m$, and the lower bound of the participation probability, denoted as q .*

The design of BEA is based on a specific mixture of information structures which helps to establish a lower bound for the regret. In particular, solving $\min_f \max_{p, g} R(f, p, g)$ is equivalent to solving a zero-sum game between nature who selects an information structure (p, g) , and a decision-maker who selects the aggregator f . We will present a specific mixed strategy of the nature, and use this mixed strategy to establish a lower bound for the regret. BEA is the best response to this mixed strategy.

LEMMA 3.2 (LOWER BOUND OF THE REGRET). *Consider a pair of information structures:*

$$^1 \alpha = \frac{(a^* q)^{n_1 - n_m}}{(a^* q)^{n_1 - n_m} + (1 - a^*)^{n_1 - n_m}}, \quad a^* \text{ is the solution to}$$

$$\max_a \sum_{s, t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-a)^{s-t}}{n((aq)^{s-t} + (1-a)^{s-t})} \right)^2$$

- $\theta_1 = (p_1 = [a, 0, \dots, 0, 1-a], g_1 = [q, 1, \dots, 1, 1])$
- $\theta_2 = (p_2 = [1-a, 0, \dots, 0, a], g_2 = [1, 1, \dots, 1, q])$

where q is the lower bound of the participation probability, and $a \in [0, 1]$ is a parameter. For any aggregator f , the regret has a lower bound

$$\begin{aligned} & \max_{p, g} R(f, p, g) \\ & \geq \max(R(f, p_1, g_1), R(f, p_2, g_2)) \\ & \geq \sum_{s, t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-a)^{s-t}}{n((aq)^{s-t} + (1-a)^{s-t})} \right)^2. \end{aligned}$$

Proof Sketch. We provide a proof sketch here, and the complete proof is deferred to Appendix A.

Consider a mixture of these two information structures: $\theta = \theta_1$ or θ_2 with equal probability. Then given the observed ratings \hat{X} , we can compute the posterior $\Pr[\theta = \theta_1 | \hat{X}]$. The best predictor for the expectation of unobserved ratings will be the conditional expectation $\mathbb{E}[x | \hat{x} = 0]$ given the information structure is θ . Then we can calculate the regret of the best predictor to obtain our lower bound.

By enumerating all possible a , we obtain a corollary

COROLLARY 3.3. *For any aggregator f ,*

$$\begin{aligned} & \max_{p, g} R(f, p, g) \\ & \geq \max_a \sum_{s, t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-a)^{s-t}}{n((aq)^{s-t} + (1-a)^{s-t})} \right)^2 \end{aligned}$$

Let $a^* = \arg \max_a \sum_{s, t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-a)^{s-t}}{n((aq)^{s-t} + (1-a)^{s-t})} \right)^2$, then according to our proof of Lemma 3.2, BEA is the best response to $\theta_1(a^*), \theta_2(a^*)$. That is, for BEA,

$$\begin{aligned} & R(f, \theta_1(a^*)) = R(f, \theta_2(a^*)) \\ & = \max_a \sum_{s, t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-\mu)}{n} \right)^2 \end{aligned}$$

Based on Corollary 3.3, BEA outperforms any other aggregator in our lower bound situation. Regarding general situations, we evaluate the performance of BEA numerically using Matlab. Figure 5 shows the results, where the grey line is the theoretical lower bound given by Corollary 3.3, the red line is BEA and the blue line is the simple averaging. BEA is almost optimal for a wide range of the parameter q , and outperforms simple averaging especially when q is small. When q is close to 1, indicating there is almost no bias, BEA is not optimal. We will discuss more details in Appendix B.

4 Rating Aggregation with Unknown Sample Size

This section will discuss the case where the number of agents n is unknown. Fortunately, when $n \rightarrow \infty$, meaning that we have a large number of agents, we can obtain a closed-form of the optimal aggregator, which is called the Polarizing-Averaging Aggregator.

When n is unknown, the aggregator only has access to the observed histogram n_1, n_2, \dots, n_m . PAA outputs the average of the empirical mean of two modified histograms. For the first histogram, we identify a threshold k_1 , and only keep q fraction of the counts for ratings above k_1 . In the second histogram, we identify a threshold k_2 , and only keep q fraction of the counts for ratings below k_2 .

Notice that in the described process, only the empirical mean of histograms is relevant. Therefore, we can use the normalized

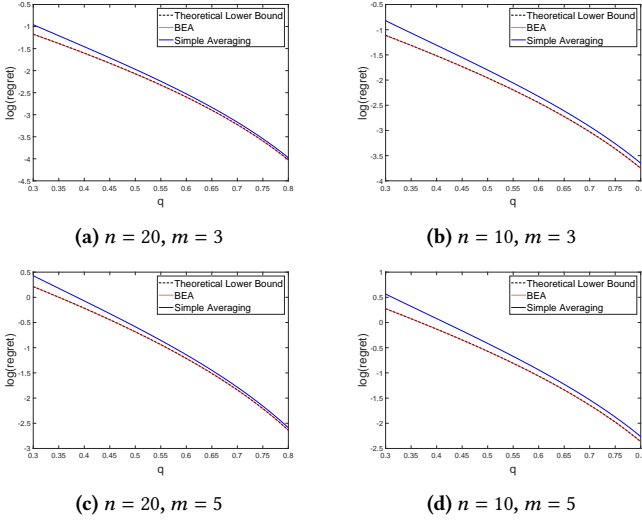


Figure 5: Simple averaging vs. BEA for different sample size n and the number of rating categories m . The x-axis is the lower bound of the participation probability, q , and the y-axis is the natural logarithm of the regret. The regret of BEA almost matches the theoretical lower bound for a wide range of q .

histograms, specifically the empirical distribution, $\hat{p}_r = \frac{n_r}{\sum_j n_j}$, $\forall r \in [m]$, as input.

DEFINITION 4.1 (POLARIZING-AVERAGING AGGREGATOR). Define the Polarizing-Averaging Aggregator $f^{\text{PAA}}(\hat{p}) = \frac{u(\hat{p}) + l(\hat{p})}{2}$, where

$$l(\hat{p}) = \frac{\sum_{r=1}^{k_1(\hat{p})} r \times \hat{p}_r + q \sum_{r=k_1(\hat{p})+1}^m r \times \hat{p}_r}{\sum_{r=1}^{k_1(\hat{p})} \hat{p}_r + q \sum_{r=k_1(\hat{p})+1}^m \hat{p}_r},$$

$$u(\hat{p}) = \frac{q \sum_{r=1}^{k_2(\hat{p})} r \times \hat{p}_r + \sum_{r=k_2(\hat{p})+1}^m r \times \hat{p}_r}{q \sum_{r=1}^{k_2(\hat{p})} \hat{p}_r + \sum_{r=k_2(\hat{p})+1}^m \hat{p}_r}.$$

We will show that $l(\hat{p})$ is the lower bound of the true expectation μ given the empirical distribution \hat{p} and the lower bound of the participation probability q , and $u(\hat{p})$ is the upper bound.

4.1 Analysis of PAA in the asymptotic case

This section will analyze PAA in the asymptotic case. We first show that in the asymptotic case, PAA is optimal. We will then present the regret and the worst information structure of PAA in the asymptotic case.

2

$$k_1(\hat{p}) = \arg \max_{1 \leq k \leq m} \sum_{r=1}^{k-1} (r-k) \hat{p}_r + q \sum_{r=k+1}^m (r-k) \hat{p}_r \geq 0$$

3

$$k_2(\hat{p}) = \arg \max_{1 \leq k \leq m} q \sum_{r=1}^{k-1} (r-k) \hat{p}_r + \sum_{r=k+1}^m (r-k) \hat{p}_r \geq 0.$$

4.1.1 Optimality of PAA.

THEOREM 4.2. When $n \rightarrow \infty$, PAA is optimal.

Proof Sketch. We provide a proof sketch here, the complete proof is deferred to Appendix A.

We first notice that when $n \rightarrow \infty$, the regret is the loss. When $n \rightarrow \infty$, the ideal aggregator who observes full ratings knows the true distribution \mathbf{p} . Therefore, the regret of f to the ideal aggregator equals the loss $\max_{\mathbf{p}, \mathbf{g}} (f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}])^2$. The problem becomes solving $\min_f \max_{\mathbf{p}, \mathbf{g}} (f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}])^2$ conditional on knowing the empirical distribution of the observed ratings $\hat{\mathbf{p}}$.

Furthermore, when $n \rightarrow \infty$, $\hat{\mathbf{p}}$ is propositional to the element-wise product⁴ of the true distribution \mathbf{p} and the participation probabilities \mathbf{g} : $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$.

The proof is divided into the following steps.

Optimal aggregator is the midpoint of extremes. Given $\hat{\mathbf{p}}$, define the lower bound of the expectation of \mathbf{p} which satisfy the propositional constraint $l^*(\hat{\mathbf{p}}) = \min_{\mathbf{p}, \mathbf{g}: \hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}} \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$ and the upper bound $u^*(\hat{\mathbf{p}}) = \max_{\mathbf{p}, \mathbf{g}: \hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}} \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$.

Since $(f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}])^2$ measures the squared distance to $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$, the best aggregator must be the midpoint between the extreme values to minimize this squared distance. That is, the best f has the format of $(u^*(\hat{\mathbf{p}}) + l^*(\hat{\mathbf{p}}))/2$.

It's left to show that $l^*(\hat{\mathbf{p}}) = l(\hat{\mathbf{p}})$ and $u^*(\hat{\mathbf{p}}) = u(\hat{\mathbf{p}})$ where $l(\hat{\mathbf{p}})$ and $u(\hat{\mathbf{p}})$ are described in the definition of PAA (Definition 4.1).

Minimize or maximize $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$ given $\hat{\mathbf{p}}$. Given $\hat{\mathbf{p}}$, we start to find $(\mathbf{p}', \mathbf{g}')$ to minimize or maximize $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$, conditional on $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$. We first characterize \mathbf{g}' . Here are two properties of the optimal \mathbf{g}' .

- \mathbf{g}' is extreme: $g_r = 1$ or q for any rating r .
- \mathbf{g}' is monotonic: $g_r \leq g_{r+1}$ or $g_r \geq g_{r+1}$ for any rating r .

Here is how we derive the above properties. Given $\hat{\mathbf{p}}$, because $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$, $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]$ can be viewed as a function, denoted as F , of \mathbf{g} . By calculating the partial derivative of F , we observe the sign of $\frac{dF}{dg_r}$ is independent with g_r as long as g_r is positive. The sign of $\frac{dF}{dg_r}$ can be determined by $F(\mathbf{g}) - r$.

Using these two properties of \mathbf{g}' , \mathbf{g}' can be constructed as follows: for the minimum, there exists an index $k_1(\hat{\mathbf{p}})$ such that $g_r = q$ for any $r \leq k_1(\hat{\mathbf{p}})$ and $g_r = 1$ otherwise. While for the maximum, there exists an index $k_2(\hat{\mathbf{p}})$ such that $g_r = 1$ for any $r \leq k_2(\hat{\mathbf{p}})$ and $g_r = q$ otherwise.

To find the optimal $k_1(\hat{\mathbf{p}})$, $k_2(\hat{\mathbf{p}})$, we can simply enumerate all the m values, but we make a more careful analysis to give the closed-form of the optimal $k_1(\hat{\mathbf{p}})$, $k_2(\hat{\mathbf{p}})$ in the appendix.

This finishes the proof sketch of the optimality of PAA in the asymptotic case.

4.1.2 Regret and the worst information structure of PAA. We also obtain the closed forms of regret and the corresponding worst information structure of PAA.

PROPOSITION 4.3. When $n \rightarrow \infty$, the regret of PAA is $\left(\frac{(m-1)(1-q)}{2(1+q)}\right)^2$. In addition, the corresponding worst pair of information structure is

- $\theta_1 = (\mathbf{p}_1 = [\frac{1}{q+1}, 0, \dots, 0, \frac{q}{q+1}], \mathbf{g}_1 = [q, 1, \dots, 1, 1])$

⁴For any vectors \mathbf{x}, \mathbf{y} , their element-wise product is $\mathbf{x} \circ \mathbf{y} = (x_1 y_1, \dots, x_n y_n)$.




Figure 6: Simple averaging vs. PAA for a specific case where the number of rating categories is $m = 3$: The x-axis is the lower bound of the participation probability, q , and the y-axis is the natural logarithm of the regret.

- $\theta_2 = (p_2 = [\frac{q}{q+1}, 0, \dots, 0, \frac{1}{q+1}], g_2 = [1, 1, \dots, 1, q])$

That is,

$$R(f^{PAA}, \theta_1) = R(f^{PAA}, \theta_2) = \left(\frac{(m-1)(1-q)}{2(1+q)} \right)^2.$$

Proof Sketch. Fix the optimal aggregator PAA, we aim to find the worst information structure (p^*, g^*) that maximizes the loss/regret $(f^{PAA}(\hat{p}) - \mathbb{E}_{x \sim p}[x])^2$.

We have proved that the optimal aggregator is the midpoint of the extremes $l(\hat{p})$ and $u(\hat{p})$. The maximal loss is $(u(\hat{p}) - l(\hat{p}))^2 / 4$. To obtain the maximal loss, we aim to find \hat{p}^* to maximize $u(\hat{p}) - l(\hat{p})$. We analyze the maximizer in the following four steps.

- We first show that \hat{p}^* has at most three non-zero entries: $|\{r : \hat{p}_r > 0\}| \leq 3$.
- We further show that \hat{p}^* has exactly two non-zero entries: $|\{r : \hat{p}_r^* > 0\}| = 2$.
- We then obtain the concrete value: $\hat{p}^* = (\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$.
- Finally, we construct two worst information structures

$$\theta_1 = (p_1 = [\frac{1}{q+1}, 0, \dots, 0, \frac{q}{q+1}], g_1 = [q, 1, \dots, 1, 1]),$$

$$\theta_2 = (p_2 = [\frac{q}{q+1}, 0, \dots, 0, \frac{1}{q+1}], g_2 = [1, 1, \dots, 1, q]).$$

In particular, θ_1 is the maximizer of $\mathbb{E}_{x \sim p}[x]$ given $\hat{p} = \hat{p}^*$.

θ_2 is the minimizer of $\mathbb{E}_{x \sim p}[x]$ given $\hat{p} = \hat{p}^*$.

First Step: For any empirical distribution \hat{p} , recall that the participation probabilities in the maximizer of $\mathbb{E}_{x \sim p}[x]$ given \hat{p} has the format that $g_r = 1$ for any $r \leq k_1(\hat{p})$ and $g_r = q$ otherwise. The participation probabilities in the minimizer of $\mathbb{E}_{x \sim p}[x]$ given \hat{p} has the format that $g_r = q$ for any $r \leq k_2(\hat{p})$ and $g_r = 1$ otherwise.

The aim is to maximize the gap $u(\hat{p}) - l(\hat{p})$. We first divide \hat{p} into three components by $k_1(\hat{p}) \leq k_2(\hat{p})$. Then we prove that by alternately concentrating all probabilities in the first component at rating 1 and all probabilities in the third component at rating m , we can achieve a new distribution with a non-decreasing gap. Note that in the new distribution, the threshold will also change. We repeat this concentration step until the distribution stabilizes at \hat{p}' , which has the following form: $(\hat{p}'_1, 0, \dots, 0, \hat{p}'_{k_1(\hat{p}')}+1, \dots, \hat{p}'_{k_2(\hat{p}')}, 0, \dots, 0, \hat{p}'_m)$. At last we concentrate all probabilities in the second component of \hat{p}' at one of rating $(k_1(\hat{p}') + 1)$ and rating $k_2(\hat{p}')$ to obtain a new

distribution with a non-decreasing gap and at most three non-zero entries.

Second Step: Now we have proved that \hat{p}^* has at most three non-zero entries at position 1, $k \in (1, m)$, m . We further find that we can redistribute the density at the middle position k to one of the end positions 1, m will not decrease the gap $u(\hat{p}) - l(\hat{p})$. Thus, we show that \hat{p}^* has exactly two non-zero entries: $|\{r : \hat{p}_r^* > 0\}| = 2$.

Third Step: Given the simple format of \hat{p}^* , we obtain its value by differentiation.

Fourth Step: Finally, we calculate the worst information structures (p^*, g^*) given \hat{p}^* by computing the maximizer and the minimizer of $\mathbb{E}_{x \sim p}[x]$ given $\hat{p} = \hat{p}^*$.

4.2 Analysis of PAA in the finite case

PAA is also applicable for the finite case. Though PAA is not optimal when n is finite, we prove that it is near-optimal, with the error bound depending on the sample size n , and the number of rating categories m .

THEOREM 4.4. When $m = o\left(\left(\frac{n}{\ln n}\right)^{\frac{1}{4}}\right)$, PAA is $O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$ -optimal.

That is, for any aggregator f ,

$$\max_{p, g} R(f^{PAA}, p, g) \leq \max_{p, g} R(f, p, g) + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right).$$

Proof Sketch. The input to f^{PAA} is the empirical distribution of observed ratings. In the asymptotic scenario, this corresponds to $p \circ g$, and we have demonstrated that f^{PAA} is optimal under these conditions. However, in finite sample situations, the empirical distribution is only an approximate version of $p \circ g$. Consequently, the performance of f^{PAA} hinges on the quality of this approximation. We will further analyze the performance by employing the Chernoff bound and union bound to assess the approximation's accuracy.

Furthermore, when n is finite, unlike the asymptotic case, the regret is not the loss. When n is different, the ideal aggregator is also different. The analysis should also consider this.

5 Virtualization of Aggregators

Visualization of BEA. Figure 7 and Figure 8 visualize BEA for a specific case where the sample size $n = 20$, and the rating values are $[m] = [2] = \{1, 2\}$. Compared to simple averaging, BEA tends to be more conservative (i.e. close to $\frac{m+1}{2} = 3/2$). Note as the lower bound of the participation probability q becomes larger, BEA becomes more conservative.

Visualization of PAA. Figure 9 and Figure 10 visualize PAA for a specific case where the sample size $n = 1000$, and the rating values are $[m] = [2] = \{1, 2\}$. Note as the lower bound of the participation probability q becomes smaller, PAA becomes more conservative (i.e. close to $\frac{m+1}{2} = 3/2$). We can see it more clearly in Figure 10.

When $n_1 = n_2 = 300$, all the aggregators report the same value $\frac{m+1}{2} = \frac{3}{2}$. In general, f^{BEA} also reports $\frac{m+1}{2}$ when $n_1 = n_m$. When n is large, the worst case is roughly $n_1 = n_m$, when $n \rightarrow \infty$, the worst \hat{p} is $(\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$. This represents a polarized situation where participants are split between extreme approval or disapproval, with equal numbers on both sides.

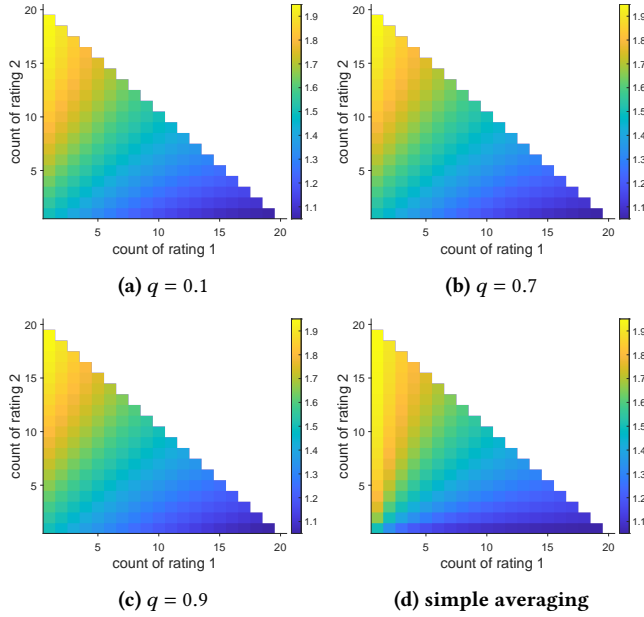


Figure 7: Heat map of BEA for a specific case where the sample size $n = 20$, and the rating values are $[m] = [2] = \{1, 2\}$. The x-axis represents the count of rating 1, and the y-axis represents the count of rating 2. We vary the lower bound of the participation probability q .




Figure 8: Value of BEA for a specific case where the sample size $n = 20$, and the rating values are $[m] = [2] = \{1, 2\}$. The count of rating 1 is fixed as $n_1 = 5$. The x-axis is the count of rating 2, n_2 . The y-axis is the value of BEA. We vary the lower bound of the participation probability q .

6 Conclusion

In this work, we explored the problem of rating aggregation with the participation. We focus on two scenarios where the sample size may be either known or unknown. For the known case, we introduced the Balanced Extremes Aggregator (BEA), which balances the extreme ratings by predicting the unobserved ratings




Figure 9: Heat map of PAA for a specific case where the sample size $n = 1000$, and the rating values are $[m] = [2] = \{1, 2\}$. The x-axis represents the count of rating 1, and the y-axis represents the count of rating 2. We vary the lower bound of the participation probability q . When $q = 1$, PAA degenerates to simple averaging.




Figure 10: Value of PAA for a specific case where the sample size $n = 1000$, and the rating values are $[m] = [2] = \{1, 2\}$. The count of rating 1 is fixed as $n_1 = 300$. The x-axis is the count of rating 2, n_2 . The y-axis is the value of PAA. We vary the lower bound of the participation probability q . When $q = 1$, PAA degenerates to simple averaging.

with the observed ratings. We evaluate its performance by numerical results. For the unknown sample size setting, we presented the Polarizing-Averaging Aggregator (PAA), which achieves near-optimal performance by averaging two polarized true distributions given the observed ratings.

As for the future work, we could explore adaptive algorithms that dynamically adjust to varying participation probabilities over time could enhance robustness, especially in rapidly changing environments like e-commerce and social media platforms.

In this work, we assume the ratings are independent. Incorporating user behavior analysis, such as identifying and adjusting for systematic biases based on user demographics or past behavior, could yield more general aggregators.

Another promising direction involves extending our aggregators to handle multi-dimensional ratings, which are common in modern applications where users rate multiple attributes (e.g., service, quality, and price).

References

- [1] Hyungtaik Ahn and James L. Powell. 1993. Semiparametric estimation of censored selection models with a nonparametric selection mechanism. *Journal of Econometrics* 58, 1-2 (1993), 3–29.
- [2] Ahu Alanya, Christof Wolf, and Cristina Sotio. 2015. Comparing multiple imputation and propensity-score weighting in unit-nonresponse adjustments: a simulation study. *Public Opinion Quarterly* 79, 3 (2015), 635–661.
- [3] Itai Arieli, Yakov Babichenko, and Rann Smorodinsky. 2018. Robust forecast aggregation. *Proceedings of the National Academy of Sciences* 115, 52 (2018), E12135–E12143.
- [4] Fabricio Benevenuto, Tiago Rodrigues, Virgílio Almeida, Jussara Almeida, and Marcos Gonçalves. 2009. Detecting spammers and content promoters in online video social networks. In *Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval*. 620–627.
- [5] Bharat Bhole and Brid Hanna. 2017. The effectiveness of online reviews in the presence of self-selection bias. *Simulation Modelling Practice and Theory* 77 (2017), 108–123.
- [6] Paul-Alexandru Chirita, Wolfgang Nejdl, and Cristian Zamfir. 2005. Preventing shilling attacks in online recommender systems. In *Proceedings of the 7th annual ACM international workshop on Web information and data management*. 67–74.
- [7] Jaime J Coon, Carena J van Riper, Lois Wright Morton, and James R Miller. 2020. Evaluating nonresponse bias in survey research conducted in the rural Midwest. *Society & Natural Resources* 33, 8 (2020), 968–986.
- [8] Simone Dahrouge, Catherine Deri Armstrong, William Hogg, Jatinderpreet Singh, and Clare Liddy. 2019. High-performing physicians are more likely to participate in a research study: findings from a quality improvement study. *BMC medical research methodology* 19 (2019), 1–9.
- [9] Henrique De Oliveira, Yuhta Ishii, and Xiao Lin. 2021. Robust merging of information. *arXiv preprint arXiv:2106.00088* (2021).
- [10] Fernando de Souza Bastos, Wagner Barreto-Souza, and Marc G Genton. 2022. A generalized Heckman model with varying sample selection bias and dispersion parameters. *Statistica Sinica* 32, 4 (2022), 1911–1938.
- [11] Dirk M Elston. 2021. Participation bias, self-selection bias, and response bias. *Journal of the American Academy of Dermatology* (2021).
- [12] Tala H Fakhouri, Crescent B Martin, Te-Ching Chen, Lara J Akinbami, Cynthia I Ogden, Ryne Paulose-Ram, Minsun K Riddles, Wendy Van de Kerckhove, Shelley B Roth, Jason Clark, et al. 2020. An investigation of nonresponse bias and survey location variability in the 2017–2018 National Health and Nutrition Examination Survey. (2020).
- [13] Quan-Yun Fu, Jian-Feng Ren, and Hong-Liang Sun. 2021. Iterative group-based and difference ranking method for online rating systems with spamming attacks. *International Journal of Modern Physics C* 32, 05 (2021), 2150059.
- [14] Ko Fujimura and Takuo Nishihara. 2003. Reputation rating system based on past behavior of evaluators. In *Proceedings of the 4th ACM Conference on Electronic Commerce*. 246–247.
- [15] Jian Gao, Yu-Wei Dong, Ming-Sheng Shang, Shi-Min Cai, and Tao Zhou. 2015. Group-based ranking method for online rating systems with spamming attacks. *Europhysics Letters* 110, 2 (2015), 28003.
- [16] Jian Gao and Tao Zhou. 2017. Evaluating user reputation in online rating systems via an iterative group-based ranking method. *Physica A: Statistical Mechanics and its Applications* 473 (2017), 546–560.
- [17] Emma Gorman, Alastair H Leyland, Gerry McCartney, Srinivasa Vittal Katikireddi, Lisa Rutherford, Lesley Graham, Mark Robinson, and Lindsay Gray. 2017. Adjustment for survey non-representativeness using record-linkage: refined estimates of alcohol consumption by deprivation in Scotland. *Addiction* 112, 7 (2017), 1270–1280.
- [18] Lindsay Gray, Emma Gorman, Ian R White, S Vittal Katikireddi, Gerry McCartney, Lisa Rutherford, and Alastair H Leyland. 2020. Correcting for non-participation bias in health surveys using record-linkage, synthetic observations and pattern mixture modelling. *Statistical methods in medical research* 29, 4 (2020), 1212–1226.
- [19] Lindsay Gray, Gerry McCartney, Ian R White, Srinivasa Vittal Katikireddi, Lisa Rutherford, Emma Gorman, and Alastair H Leyland. 2013. Use of record-linkage to handle non-response and improve alcohol consumption estimates in health survey data: a study protocol. *BMJ open* 3, 3 (2013), e002647.
- [20] Yongkang Guo, Jason D. Hartline, Zhihuan Huang, Yuqing Kong, Anant Shah, and Fang-Yi Yu. 2024. Algorithmic Robust Forecast Aggregation. *arXiv:2401.17743* [cs.LG].
- [21] James Heckman. 1974. Shadow prices, market wages, and labor supply. *Econometrica: Journal of the econometric society* (1974), 679–694.
- [22] J Heckman. 1979. Sample selection bias as a specification error. *Econometrica* (1979).
- [23] James J Heckman. 1976. The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. In *Annals of economic and social measurement, volume 5, number 4*. NBER, 475–492.
- [24] James J Heckman and Richard Robb Jr. 1985. Alternative methods for evaluating the impact of interventions: An overview. *Journal of econometrics* 30, 1-2 (1985), 239–267.
- [25] Nan Hu, Paul A Pavlou, and Jie Jennifer Zhang. 2009. Why do online product reviews have a J-shaped distribution? Overcoming biases in online word-of-mouth communication. *Commun. ACM* 52, 10 (2009), 144–147.
- [26] Audun Jøsang, Roslan Ismail, and Colin Boyd. 2007. A survey of trust and reputation systems for online service provision. *Decision support systems* 43, 2 (2007), 618–644.
- [27] Noi Sian Koh, Nan Hu, and Eric K Clemons. 2010. Do online reviews reflect a product's true perceived quality? An investigation of online movie reviews across cultures. *Electronic commerce research and applications* 9, 5 (2010), 374–385.
- [28] Yuqing Kong, Shu Wang, and Ying Wang. 2024. The Surprising Benefits of Base Rate Neglect in Robust Aggregation. *arXiv preprint arXiv:2406.13490* (2024).
- [29] Paolo Laureti, Lionel Moret, Y-C Zhang, and Y-K Yu. 2006. Information filtering via iterative refinement. *Europhysics Letters* 75, 6 (2006), 1006.
- [30] Lung-Fei Lee. 1982. Some approaches to the correction of selectivity bias. *The Review of Economic Studies* 49, 3 (1982), 355–372.
- [31] Lung-Fei Lee. 1983. Generalized econometric models with selectivity. *Econometrica: Journal of the Econometric Society* (1983), 507–512.
- [32] Lung-Fei Lee. 1994. Semiparametric instrumental variable estimation of simultaneous equation sample selection models. *Journal of Econometrics* 63, 2 (1994), 341–388.
- [33] Hao Liao, An Zeng, Rui Xiao, Zhuo-Ming Ren, Duan-Bing Chen, and Yi-Cheng Zhang. 2014. Ranking reputation and quality in online rating systems. *PloS one* 9, 5 (2014), e97146.
- [34] Roderick JA Little. 1993. Pattern-mixture models for multivariate incomplete data. *J. Amer. Statist. Assoc.* 88, 421 (1993), 125–134.
- [35] Yulia V Marchenko and Marc G Genton. 2012. A Heckman selection-t model. *J. Amer. Statist. Assoc.* 107, 497 (2012), 304–317.
- [36] Arjun Mukherjee, Bing Liu, and Natalie Glance. 2012. Spotting fake reviewer groups in consumer reviews. In *Proceedings of the 21st international conference on World Wide Web*. 191–200.
- [37] Eric Neyman and Tim Roughgarden. 2022. Are you smarter than a random expert? The robust aggregation of substitutable signals. In *Proceedings of the 23rd ACM Conference on Economics and Computation*. 990–1012.
- [38] Yuqi Pan, Zhaohua Chen, and Yuqing Kong. 2023. Robust Decision Aggregation with Second-order Information. *arXiv preprint arXiv:2311.14094* (2023).
- [39] Andy Peytchev. 2012. Multiple imputation for unit nonresponse and measurement error. *Public Opinion Quarterly* 76, 2 (2012), 214–237.
- [40] Paul Resnick, Ko Kuwabara, Richard Zeckhauser, and Eric Friedman. 2000. Reputation systems. *Commun. ACM* 43, 12 (2000), 45–48.
- [41] James M Robins, Andrea Rotnitzky, and Lue Ping Zhao. 1994. Estimation of regression coefficients when some regressors are not always observed. *Journal of the American statistical Association* 89, 427 (1994), 846–866.
- [42] Tabea Schoeler, Doug Speed, Eleonora Porcu, Nicola Pirastu, Jean-Baptiste Pingault, and Zoltán Kutalik. 2022. Correction for participation bias in the UK Biobank reveals non-negligible impact on genetic associations and downstream analyses. *bioRxiv* (2022), 2022–09.
- [43] Tabea Schoeler, Doug Speed, Eleonora Porcu, Nicola Pirastu, Jean-Baptiste Pingault, and Zoltán Kutalik. 2023. Participation bias in the UK Biobank distorts genetic associations and downstream analyses. *Nature Human Behaviour* 7, 7 (2023), 1216–1227.
- [44] Jonathan AC Sterne, Ian R White, John B Carlin, Michael Spratt, Patrick Royston, Michael G Kenward, Angela M Wood, and James R Carpenter. 2009. Multiple imputation for missing data in epidemiological and clinical research: potential and pitfalls. *Bmj* 338 (2009).
- [45] Armin Tabandeh, Gaofeng Jia, and Paolo Gardoni. 2022. A review and assessment of importance sampling methods for reliability analysis. *Structural Safety* 97 (2022), 102216.
- [46] Surya T Tokdar and Robert E Kass. 2010. Importance sampling: a review. *Wiley Interdisciplinary Reviews: Computational Statistics* 2, 1 (2010), 54–60.

- [47] Daniel Mark Tompsett, Finbarr Leacy, Margarita Moreno-Betancur, Jon Heron, and Ian R White. 2018. On the use of the not-at-random fully conditional specification (NARFCS) procedure in practice. *Statistics in medicine* 37, 15 (2018), 2338–2353.
- [48] Pierre Verger, Dimitri Scronias, Yves Fradier, Malika Meziani, and Bruno Ventelou. 2021. Online study of health professionals about their vaccination attitudes and behavior in the COVID-19 era: addressing participation bias. *Human Vaccines & Immunotherapeutics* 17, 9 (2021), 2934–2939.
- [49] Robin Vogel, Mastane Achab, Stéphan Cléménçon, and Charles Tillier. 2020. Weighted empirical risk minimization: Sample selection bias correction based on importance sampling. *arXiv preprint arXiv:2002.05145* (2020).
- [50] Zhuo Wang, Songmin Gu, Xiangnan Zhao, and Xiaowei Xu. 2018. Graph-based review spammer group detection. *Knowledge and Information Systems* 55 (2018), 571–597.
- [51] Zhuo Wang, Tingting Hou, Dawei Song, Zhun Li, and Tianqi Kong. 2016. Detecting review spammer groups via bipartite graph projection. *Comput. J.* 59, 6 (2016), 861–874.
- [52] Fuzhi Zhang, Xiaoyan Hao, Jinbo Chao, and Shuai Yuan. 2020. Label propagation-based approach for detecting review spammer groups on e-commerce websites. *Knowledge-Based Systems* 193 (2020), 105520.
- [53] Mikhail Zhelonkin, Marc G Genton, and Elvezio Ronchetti. 2016. Robust inference in sample selection models. *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78, 4 (2016), 805–827.
- [54] Yan-Bo Zhou, Ting Lei, and Tao Zhou. 2011. A robust ranking algorithm to spamming. *Europhysics Letters* 94, 4 (2011), 48002.
- [55] Huan Zhu, Yu Xiao, Dongmei Chen, and Jun Wu. 2024. A Robust Rating Aggregation Method based on Rater Group Trustworthiness for Collusive Disturbance. *Information Systems Frontiers* (2024), 1–21.
- [56] Huan Zhu, Yu Xiao, Zhi-Gang Wang, and Jun Wu. 2023. A robust reputation iterative algorithm based on Z-statistics in a rating system with thorny objects. *Journal of the Operational Research Society* 74, 6 (2023), 1600–1612.
- [57] Qian Zhu, Leo Yu-Ho Lo, Meng Xia, Zixin Chen, and Xiaojuan Ma. 2022. Bias-aware design for informed decisions: Raising awareness of self-selection bias in user ratings and reviews. *Proceedings of the ACM on Human-Computer Interaction* 6, CSCW2 (2022), 1–31.

A Omitted Proofs

LEMMA A.1 (LOWER BOUND OF THE REGRET). *Consider a pair of information structures:*

- $\theta_1 = (\mathbf{p}_1 = [a, 0, \dots, 0, 1 - a], \mathbf{g}_1 = [q, 1, \dots, 1, 1])$
- $\theta_2 = (\mathbf{p}_2 = [1 - a, 0, \dots, 0, a], \mathbf{g}_2 = [1, 1, \dots, 1, q])$

where a is a parameter. Then for any aggregator f , the regret has a lower bound

$$\begin{aligned} & \max_{\mathbf{p}, \mathbf{g}} R(f, \mathbf{p}, \mathbf{g}) \\ & \geq \max(R(f, \mathbf{p}_1, \mathbf{g}_1), R(f, \mathbf{p}_2, \mathbf{g}_2)) \\ & \geq \sum_{s,t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-a)^{s-t}}{n((aq)^{s-t} + (1-a)^{s-t})} \right)^2. \end{aligned}$$

PROOF. Consider a linear combination information structure $\theta = \theta_1$ or θ_2 with equal probability. Now we compute its best response $f'(\hat{\mathbf{X}}) = \mathbb{E}[\frac{1}{n} \sum_i x_i | \hat{\mathbf{X}}]$, which is the posterior given the observed ratings. Given $s = \sum_i \mathbb{1}(\hat{x}_i = 1)$, $t = \sum_i \mathbb{1}(\hat{x}_i = m)$. We have

$$\begin{aligned} & \Pr[\theta = \theta_1 | \hat{\mathbf{X}}] \\ & = \frac{\Pr[\hat{\mathbf{X}} | \theta_1]}{\Pr[\hat{\mathbf{X}} | \theta_1] + \Pr[\hat{\mathbf{X}} | \theta_2]} \\ & = \frac{\binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t}}{\binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} + \binom{n}{t} \binom{n-t}{s} a^{n-s} (1-a)^s q^t (1-q)^{n-s-t}} \\ & = \frac{(aq)^{s-t}}{(aq)^{s-t} + (1-a)^{s-t}}. \\ & = \frac{1}{1 + (\frac{1-a}{aq})^{s-t}}. \end{aligned}$$

Then

$$\begin{aligned} & \mathbb{E}[x | \hat{\mathbf{X}} = 0] \\ & = \Pr[\theta = \theta_1 | \hat{\mathbf{X}}] \mathbb{E}[x | \hat{\mathbf{X}} = 0, \theta_1] + \Pr[\theta = \theta_2 | \hat{\mathbf{X}}] \mathbb{E}[x | \hat{\mathbf{X}} = 0, \theta_2] \\ & = 1 * \frac{1}{1 + (\frac{1-a}{aq})^{s-t}} + m * \left(1 - \frac{1}{1 + (\frac{1-a}{aq})^{s-t}} \right) \\ & = m - (m-1) \frac{1}{1 + (\frac{1-a}{aq})^{s-t}}. \end{aligned}$$

Thus the best response is $f'(\hat{\mathbf{X}}) = \frac{\sum_i \hat{x}_i + (m-1) \frac{1}{1 + (\frac{1-a}{aq})^{s-t}} * (n-s-t)}{n}$.

By simple calculation we obtain that for any f ,

$$R(f, \Theta) \geq R(f, \theta) \geq R(f', \theta)$$

$$= \sum_{s,t} \binom{n}{t} \binom{n-t}{s} a^{n-t} (1-a)^t q^s (1-q)^{n-s-t} \left(\frac{(n-s-t)(m-1)(1-\mu)}{n} \right)^2$$

□

THEOREM A.2. *When $n \rightarrow \infty$, the aggregator is optimal.*

PROOF. When $n \rightarrow \infty$, the empirical distribution is propositional to the element-wise product of the true distribution \mathbf{p} and the participation probabilities \mathbf{g} : $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$. We want to minimize $\max_{\mathbf{p}, \mathbf{g}} (f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x])^2$.

Given $\hat{\mathbf{p}}$, define the lower bound of the expectation of \mathbf{p} which satisfy the propositional constraint $l^*(\hat{\mathbf{p}}) = \min_{\mathbf{p}, \mathbf{g}: \hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}} \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x]$ and the upper bound $u^*(\hat{\mathbf{p}}) = \max_{\mathbf{p}, \mathbf{g}: \hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}} \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x]$. Since $(f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x])^2$ is a quadratic function about $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x]$, we have

$$\begin{aligned} \max_{\mathbf{p}, \mathbf{g}} (f(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x])^2 &= \max\{(f(\hat{\mathbf{p}}) - l^*(\hat{\mathbf{p}}))^2, (f(\hat{\mathbf{p}}) - u^*(\hat{\mathbf{p}}))^2\} \\ &\geq (u^*(\hat{\mathbf{p}}) - l^*(\hat{\mathbf{p}}))^2 / 4. \end{aligned}$$

The equality holds if and only if $f(\hat{\mathbf{p}}) = (u^*(\hat{\mathbf{p}}) + l^*(\hat{\mathbf{p}})) / 2$, so the best aggregator is $f(\hat{\mathbf{p}}) = (u^*(\hat{\mathbf{p}}) + l^*(\hat{\mathbf{p}})) / 2$. It's left to show that $l^*(\hat{\mathbf{p}}) = l(\hat{\mathbf{p}})$ and $u^*(\hat{\mathbf{p}}) = u(\hat{\mathbf{p}})$ where $l(\hat{\mathbf{p}})$ and $u(\hat{\mathbf{p}})$ are described in the definition of PAA (Definition 4.1).

Given $\hat{\mathbf{p}}$, we start to find $(\mathbf{p}', \mathbf{g}')$ to minimize or maximize $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x]$, conditional on $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$. We first characterize \mathbf{g}' . Here are two properties of the optimal \mathbf{g}' .

- \mathbf{g}' is extreme: $g_r = 1$ or q for any rating r .
- \mathbf{g}' is monotonic: $g_r \leq g_{r+1}$ or $g_r \geq g_{r+1}$ for any rating r .

Given $\hat{\mathbf{p}}$, because $\hat{\mathbf{p}} \propto \mathbf{p} \circ \mathbf{g}$, $\mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x]$ can be viewed as a function, denoted as F , of \mathbf{g} . By calculating the partial derivative of F , we have

$$\begin{aligned} \frac{dF}{dg_r} &= \frac{pr \hat{g}_r^{-2}}{(\sum_j \hat{p}_j / g_j)^2} * (\sum_j \frac{\hat{p}_j}{g_j} - r \sum_j \frac{\hat{p}_j}{g_j}) \\ &= \frac{pr \hat{g}_r^{-2}}{(\sum_j \hat{p}_j / g_j)^2} * (\sum_{j=1}^{r-1} (j-r) \frac{\hat{p}_j}{g_j} + \sum_{j=r+1}^m (j-r) \frac{\hat{p}_j}{g_j}) \\ &= \frac{pr \hat{g}_r^{-2}}{\sum_j \hat{p}_j / g_j} * (F(\mathbf{g}) - r). \end{aligned}$$

First Notice that the sign of $\frac{dF}{dg_r}$ is independent with g_r as long as g_r is positive, so the optimal \mathbf{g} is extreme: $g_r = 1$ or $g_r = q$. Then Notice the sign of $\frac{dF}{dg_r}$ is determined by $F(\mathbf{g}) - r$ where $F(\mathbf{g}) = \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[x] \in [1, m]$, so the optimal \mathbf{g} is monotonic.

Using these two properties of g', g' can be given by an index strategy:

for the minimum, there exists an index $k_1(\hat{p})$ such that

$$g'_r = \begin{cases} q & r \leq k_1(\hat{p}) \\ 1 & r > k_1(\hat{p}) \end{cases}$$

while for the maximum, there exists an index $k_2(\hat{p})$ such that

$$g'_r = \begin{cases} 1 & r \leq k_2(\hat{p}) \\ q & r > k_2(\hat{p}) \end{cases}$$

To find the optimal index, we can simply calculate all the m values, but here we make a more careful analysis.

Take the minimum for example. We use $g^{(k)}$ to denote a $1 \times m$ vector with the first k elements being q and the other elements being 1. We claim $F(g^{(k)})$ has a single-bottom shape regarding the index k , or equivalently,

$$F(g^{(k)}) - F(g^{(k-1)}) > 0 \implies \forall k' \geq k, F(g^{(k')}) - F(g^{(k'-1)}) \geq 0.$$

Let $D_k = \frac{1}{q} \sum_{i=1}^k \hat{p}_i + \sum_{i=k+1}^m \hat{p}_i$. We have

$$F(g^{(k)}) - F(g^{(k-1)}) = \left(\frac{1}{q} - 1\right) \hat{p}_k D_k (k - F(g^{(k)})).$$

We first prove $F(g^{(k)}) < k \implies \forall k' \geq k, F(g^{(k')}) < k'$. Notice

$$F(g^{(k+1)}) - F(g^{(k)}) = \left(\frac{1}{q} - 1\right) \hat{p}_{k+1} D_{k+1} ((k+1) - F(g^{(k+1)})).$$

So

$$F(g^{(k+1)}) = \frac{F(k) + \left(\frac{1}{q} - 1\right)(k+1)\hat{p}_{k+1}D_{k+1}}{1 + \left(\frac{1}{q} - 1\right)\hat{p}_{k+1}D_{k+1}}.$$

If $F(g^{(k)}) < k$, we have

$$\begin{aligned} F(g^{(k+1)}) &= \frac{F(k) + \left(\frac{1}{q} - 1\right)(k+1)\hat{p}_{k+1}D_{k+1}}{1 + \left(\frac{1}{q} - 1\right)\hat{p}_{k+1}D_{k+1}} \\ &< \frac{k + \left(\frac{1}{q} - 1\right)(k+1)\hat{p}_{k+1}D_{k+1}}{1 + \left(\frac{1}{q} - 1\right)\hat{p}_{k+1}D_{k+1}} \\ &< \frac{(k+1) + \left(\frac{1}{q} - 1\right)(k+1)\hat{p}_{k+1}D_{k+1}}{1 + \left(\frac{1}{q} - 1\right)\hat{p}_{k+1}D_{k+1}} \\ &= k+1 \end{aligned}$$

By induction, we have $\forall k' \geq k, F(g^{(k')}) < k'$. So

$$\begin{aligned} F(g^{(k)}) - F(g^{(k-1)}) > 0 &\implies F(g^{(k)}) < k \\ &\implies \forall k' \geq k, F(g^{(k')}) < k' \\ &\implies \forall k' \geq k, F(g^{(k')}) - F(g^{(k'-1)}) \geq 0. \end{aligned}$$

So $k_1(\hat{p})$ is the optimal index for minimum $\iff F(g^{(k_1(\hat{p}))}) - k_1(\hat{p}) > 0, F(g^{(k_1(\hat{p})+1)}) - (k_1(\hat{p}) + 1) < 0$ (Since \hat{p} can have zero entries, there might be some consecutive k that all of them is "optimal". For clarity, we take this definition). i.e.

$$k_1(\hat{p}) = \max\{k : 1 \leq k \leq m, \sum_{i=1}^{k-1} (i-k) \frac{\hat{p}_i}{q} + \sum_{i=k}^m (i-k) \hat{p}_i \geq 0\}.$$

Similarly,

$$k_2(\hat{p}) = \max\{k : 1 \leq k \leq m, \sum_{i=1}^{k-1} (i-k) \hat{p}_i + \sum_{i=k}^m (i-k) \frac{\hat{p}_i}{q} \geq 0\}.$$

Notice the optimal index is also a rough approximation of the optimal value. Since $k_1(\hat{p}) < l(\hat{p}) < k_1(\hat{p}) + 1, k_2(\hat{p}) < u(\hat{p}) < k_2(\hat{p}) + 1$ and $l(\hat{p}) < u(\hat{p})$, we have $k_1(\hat{p}) \leq k_2(\hat{p})$. \square

PROPOSITION A.3. When $n \rightarrow \infty$, the regret of PAA is $\left(\frac{(m-1)(1-q)}{2(1+q)}\right)^2$. In addition, the corresponding worst pair of information structure is

- $\theta_1 = (p_1 = [\frac{1}{q+1}, 0, \dots, 0, \frac{q}{q+1}], g_1 = [q, 1, \dots, 1, 1])$
- $\theta_2 = (p_2 = [\frac{q}{q+1}, 0, \dots, 0, \frac{1}{q+1}], g_2 = [1, 1, \dots, 1, q])$

That is,

$$R(f^{PAA}, \theta_1) = R(f^{PAA}, \theta_2) = \left(\frac{(m-1)(1-q)}{2(1+q)}\right)^2.$$

PROOF. Fix the optimal aggregator PAA, we aim to find the worst information structure (p^*, g^*) that maximizes the loss/regret $(f^{PAA}(\hat{p}) - \mathbb{E}_{x \sim p}[x])^2$.

We have proved that the optimal aggregator is the midpoint of the extremes $l(\hat{p})$ and $u(\hat{p})$. The maximal loss is $(u(\hat{p}) - l(\hat{p}))^2 / 4$. To obtain the maximal loss, we aim to find \hat{p}^* to maximize $u(\hat{p}) - l(\hat{p})$. We analyze the maximizer in the following four steps.

First we prove \hat{p}^* has at most three non-zero entries. Then we prove \hat{p}^* has exactly two non-zero entries. Next we prove the \hat{p}^* is $(\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$. Finally, we calculate the worst information structures (p^*, g^*) given \hat{p}^* by computing the maximizer and the minimizer of $\mathbb{E}_{x \sim p}[x]$ given $\hat{p} = \hat{p}^*$.

First Step: For any \hat{p} , consider the optimal index for minimum $k_1(\hat{p})$ and the optimal index for maximum $k_2(\hat{p})$. Since $k_1(\hat{p}) < l(\hat{p}) < k_1(\hat{p}) + 1, k_2(\hat{p}) < u(\hat{p}) < k_2(\hat{p}) + 1$ and $l(\hat{p}) < u(\hat{p})$, we have $k_1(\hat{p}) \leq k_2(\hat{p})$. Define $a = \sum_{j=1}^{k_1(\hat{p})} \hat{p}_j$ and

$$\hat{p}_1 = (a, 0, \dots, 0, \hat{p}_{k_1(\hat{p})+1}, \hat{p}_{k_1(\hat{p})+2}, \dots, \hat{p}_m).$$

For short, we denote $R(f, \hat{p}) = \max_{p, g: \hat{p} \in p \circ g} (f(\hat{p}) - \mathbb{E}_{x \sim p}[x])^2$, which is the maximum regret of f given the empirical distribution \hat{p} . We will show $R(f^{PAA}, \hat{p}) \leq R(f^{PAA}, \hat{p}_1)$.

Note for \hat{p}_1 , we have

$$l(\hat{p}_1) \leq \frac{\frac{1}{q} \sum_{j=1}^{k_1(\hat{p})} \hat{p}_j + \sum_{j=k_1(\hat{p})+1}^m j \hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1(\hat{p})} \hat{p}_j + \sum_{j=k_1(\hat{p})+1}^m \hat{p}_j},$$

and

$$u(\hat{p}_1) \geq \frac{\sum_{j=1}^{k_2(\hat{p})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{p})+1}^m j \hat{p}_j}{\sum_{j=1}^{k_2(\hat{p})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{p})+1}^m \hat{p}_j}.$$

So

$$\begin{aligned}
& (u(\hat{\mathbf{p}}_1) - l(\hat{\mathbf{p}}_1)) - (u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}})) \\
&= (l(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}}_1)) - (u(\hat{\mathbf{p}}) - u(\hat{\mathbf{p}}_1)) \\
&\geq \left(\frac{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} j\hat{p}_j + \sum_{j=k_1(\hat{\mathbf{p}})+1}^m j\hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} \hat{p}_j + \sum_{j=k_1(\hat{\mathbf{p}})+1}^m \hat{p}_j} - \frac{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} j\hat{p}_j + \sum_{j=k_1(\hat{\mathbf{p}})+1}^m j\hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} \hat{p}_j + \sum_{j=k_1(\hat{\mathbf{p}})+1}^m \hat{p}_j} \right) \\
&\quad - \left(\frac{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} j\hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m j\hat{p}_j}{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m \hat{p}_j} - \frac{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} j\hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m j\hat{p}_j}{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m \hat{p}_j} \right) \\
&= \frac{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} (j-1)\hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}})} \hat{p}_j + \sum_{j=k_1(\hat{\mathbf{p}})+1}^m \hat{p}_j} - \frac{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} (j-1)\hat{p}_j}{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m \hat{p}_j} \\
&= \frac{\sum_{j=1}^{k_1(\hat{\mathbf{p}})} (j-1)\hat{p}_j}{\sum_{j=1}^{k_1(\hat{\mathbf{p}})} \hat{p}_j + q \sum_{j=k_1(\hat{\mathbf{p}})+1}^m \hat{p}_j} - \frac{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} (j-1)\hat{p}_j}{\sum_{j=1}^{k_2(\hat{\mathbf{p}})} \hat{p}_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}})+1}^m \hat{p}_j} \\
&\geq 0.
\end{aligned}$$

which implies $R(f^{PAA}, \hat{\mathbf{p}}) \leq R(f^{PAA}, \hat{\mathbf{p}}_1)$.

We also use $\mathbf{p}(1)$ to denote the first element of \mathbf{p} . Define $b = \sum_{j=k_2(\hat{\mathbf{p}})+1}^m \hat{\mathbf{p}}_1(j)$ and $\hat{\mathbf{p}}_2 = (\hat{\mathbf{p}}_1(1), \hat{\mathbf{p}}_1(2), \dots, \hat{\mathbf{p}}_1(k_2(\hat{\mathbf{p}})), 0, \dots, 0, b)$.

Similarly, we can prove $R(f^{PAA}, \hat{\mathbf{p}}_1) \leq R(f^{PAA}, \hat{\mathbf{p}}_2)$.

We call this one iteration. Note the non-zero entries will not increase during the iteration. We iterate until the non-zero entries do not decrease. We use $\hat{\mathbf{p}}'$ to denote the stable distribution after iterations. $\hat{\mathbf{p}}'$ has the format of $(\hat{p}'_1, 0, \dots, 0, \hat{p}'_l, \hat{p}'_{l+1}, \dots, \hat{p}'_r, 0, \dots, 0, \hat{p}'_m)$, where $k_1(\hat{\mathbf{p}}') < l$ and $k_2(\hat{\mathbf{p}}') > r$.

Define $\hat{\mathbf{p}}_3 = (\hat{p}'_1, 0, \dots, 0, 1 - \hat{p}'_1 - \hat{p}'_m, 0, \dots, 0, \hat{p}'_m)$ where the l -th entry is $1 - \hat{p}'_1 - \hat{p}'_m$. Define $\hat{\mathbf{p}}_4 = (\hat{p}'_1, 0, \dots, 0, 1 - \hat{p}'_1 - \hat{p}'_m, 0, \dots, 0, \hat{p}'_m)$ where the r -th entry is $1 - \hat{p}'_1 - \hat{p}'_m$.

$$l(\hat{\mathbf{p}}_3) \leq \frac{\frac{1}{q} \sum_{j=1}^{l-1} j\hat{p}_3(j) + \sum_{j=l}^m j\hat{p}_3(j)}{\frac{1}{q} \sum_{j=1}^{l-1} \hat{p}_3(j) + \sum_{j=l}^m \hat{p}_3(j)} = \frac{\frac{1}{q} \hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + m\hat{p}'_m}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)}$$

$$u(\hat{\mathbf{p}}_3) \geq \frac{\sum_{j=1}^r j\hat{p}_3(j) + \frac{1}{q} \sum_{j=r+1}^m j\hat{p}_3(j)}{\sum_{j=1}^r \hat{p}_3(j) + \frac{1}{q} \sum_{j=r+1}^m \hat{p}_3(j)} = \frac{\hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + \frac{m}{q} \hat{p}'_m}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m}$$

If $\hat{p}'_1 \leq \hat{p}'_m$, then

$$\begin{aligned}
& (u(\hat{\mathbf{p}}_3) - l(\hat{\mathbf{p}}_3)) - (u(\hat{\mathbf{p}}') - l(\hat{\mathbf{p}}')) \\
&= (l(\hat{\mathbf{p}}') - l(\hat{\mathbf{p}}_3)) - (u(\hat{\mathbf{p}}') - u(\hat{\mathbf{p}}_3)) \\
&\geq \left(\frac{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}}')} j\hat{p}'_j + \sum_{j=k_1(\hat{\mathbf{p}}')+1}^m j\hat{p}'_j}{\frac{1}{q} \sum_{j=1}^{k_1(\hat{\mathbf{p}}')} \hat{p}'_j + \sum_{j=k_1(\hat{\mathbf{p}}')+1}^m \hat{p}'_j} - \frac{\frac{1}{q} \hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + m\hat{p}'_m}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)} \right) \\
&\quad - \left(\frac{\sum_{j=1}^{k_2(\hat{\mathbf{p}}')} j\hat{p}'_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}}')+1}^m j\hat{p}'_j}{\sum_{j=1}^{k_2(\hat{\mathbf{p}}')} \hat{p}'_j + \frac{1}{q} \sum_{j=k_2(\hat{\mathbf{p}}')+1}^m \hat{p}'_j} - \frac{\hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + \frac{m}{q} \hat{p}'_m}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m} \right) \\
&= \left(\frac{\frac{1}{q} \hat{p}'_1 + \sum_{j=l}^r j\hat{p}'_j + m\hat{p}'_m}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)} - \frac{\frac{1}{q} \hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + m\hat{p}'_m}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)} \right) \\
&\quad - \left(\frac{\hat{p}'_1 + \sum_{j=l}^r j\hat{p}'_j + \frac{m}{q} \hat{p}'_m}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m} - \frac{\hat{p}'_1 + l(1 - \hat{p}'_1 - \hat{p}'_m) + \frac{m}{q} \hat{p}'_m}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m} \right) \\
&= \left(\sum_{j=l}^r j\hat{p}'_j - l(1 - \hat{p}'_1 - \hat{p}'_m) \right) \left(\frac{1}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)} - \frac{1}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m} \right) \\
&= \sum_{j=l}^r (j-l)\hat{p}'_j \left(\frac{1}{\frac{1}{q} \hat{p}'_1 + (1 - \hat{p}'_1)} - \frac{1}{(1 - \hat{p}'_m) + \frac{1}{q} \hat{p}'_m} \right) \\
&\geq 0.
\end{aligned}$$

Similarly, if $\hat{p}'_1 > \hat{p}'_m$, we have

$$(u(\hat{\mathbf{p}}_4) - l(\hat{\mathbf{p}}_4)) - (u(\hat{\mathbf{p}}') - l(\hat{\mathbf{p}}')) \geq 0.$$

So $(u(\hat{\mathbf{p}}') - l(\hat{\mathbf{p}}')) \leq \max\{(u(\hat{\mathbf{p}}_3) - l(\hat{\mathbf{p}}_3)), (u(\hat{\mathbf{p}}_4) - l(\hat{\mathbf{p}}_4))\}$, which implies $R(f^{PAA}, \hat{\mathbf{p}}') \leq \max\{R(f^{PAA}, \hat{\mathbf{p}}_3), R(f^{PAA}, \hat{\mathbf{p}}_4)\}$. So $\hat{\mathbf{p}}^*$ has at most three non-zero entries.

Second Step: Now we prove $\hat{\mathbf{p}}^*$ has exactly two non-zero entries. Without loss of generality, we now suppose there are only three ratings: 1, k , m , where $1 < k < m$. For $\hat{\mathbf{p}} = (a, 1 - a - b, b)$, suppose the optimal index for minimum is $k_1(\hat{\mathbf{p}}) \in \{1, 2\}$ and the optimal index for maximum is $k_2(\hat{\mathbf{p}}) \in \{1, 2\}$. Still we have $k_1(\hat{\mathbf{p}}) \leq k_2(\hat{\mathbf{p}})$. Define $\hat{\mathbf{p}}_5 = (a, 0, 1 - a)$, $\hat{\mathbf{p}}_6 = (1 - b, 0, b)$. There are three cases.

Case 1: We have $k_1(\hat{\mathbf{p}}) = k_2(\hat{\mathbf{p}}) = 1$.

$$l(\hat{\mathbf{p}}) = \frac{a}{a + q(1-a)} + \frac{q(1-a-b)}{a + q(1-a)} \times k + \frac{qb}{a + q(1-a)} \times m$$

$$u(\hat{\mathbf{p}}) = \frac{qa}{qa + 1 - a} + \frac{1 - a - b}{qa + (1-a)} \times k + \frac{b}{qa + (1-a)} \times m$$

$$l(\hat{\mathbf{p}}_5) = \frac{a}{a + q(1-a)} + \frac{q(1-a)}{a + q(1-a)} \times m$$

$$u(\hat{\mathbf{p}}_5) = \frac{qa}{qa + (1-a)} + \frac{1-a}{qa + (1-a)} \times m$$

Then

$$\begin{aligned}
& (u(\hat{\mathbf{p}}_5) - l(\hat{\mathbf{p}}_5)) - (u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}})) \\
&= (u(\hat{\mathbf{p}}_5) - u(\hat{\mathbf{p}})) - (l(\hat{\mathbf{p}}_5) - l(\hat{\mathbf{p}})) \\
&= \frac{1-a-b}{qa + (1-a)} \times k - \frac{q(1-a-b)}{a + q(1-a)} \times k \\
&= k(1-a-b) \left(\frac{1}{qa + (1-a)} - \frac{1}{\frac{1}{q} + (1-a)} \right) \\
&\geq 0.
\end{aligned}$$

So $R(f^{PAA}, \hat{\mathbf{p}}) \leq R(f^{PAA}, \hat{\mathbf{p}}^5)$.

Case 2: $k_1(\hat{\mathbf{p}}) = k_2(\hat{\mathbf{p}}) = 2$. Similarly, we can prove $R(f^{PAA}, \hat{\mathbf{p}}) \leq R(f^{PAA}, \hat{\mathbf{p}}_6)$.

Case 3: $k_1(\hat{\mathbf{p}}) = 1, k_2(\hat{\mathbf{p}}) = 2$.

$$l(\hat{\mathbf{p}}) = \frac{a}{a+q(1-a)} + \frac{q(1-a-b)}{a+q(1-a)} \times k + \frac{qb}{a+q(1-a)} \times m$$

$$u(\hat{\mathbf{p}}) = \frac{qa}{b+q(1-b)} + \frac{q(1-a-b)}{b+q(1-b)} \times k + \frac{b}{b+q(1-b)} \times m$$

Define

$$A = \frac{q(1-b)}{b+q(1-b)} + \frac{bm}{b+q(1-b)} - (1 + \frac{qb(m-1)}{a+q(1-a)}),$$

$$B = 1 + \frac{(q(1-a-b)+b)(m-1)}{b+q(1-b)} - (\frac{a}{a+q(1-a)} + \frac{q(1-a)m}{a+q(1-a)}).$$

Since $u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}})$ is linear on k , we have

$$u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}}) \leq \max\{A, B\}.$$

Notice

$$u(\hat{\mathbf{p}}_6) = \frac{q(1-b)}{b+q(1-b)} + \frac{bm}{b+q(1-b)}, \quad l(\hat{\mathbf{p}}_5) = \frac{a}{a+q(1-a)} + \frac{q(1-a)m}{a+q(1-a)}.$$

So

$$A = u(\hat{\mathbf{p}}_6) - (1 + \frac{qb(m-1)}{a+q(1-a)}), \quad B = 1 + \frac{(q(1-a-b)+b)(m-1)}{b+q(1-b)} - l(\hat{\mathbf{p}}_5)$$

$$\begin{aligned} & l(\hat{\mathbf{p}}_6) - (1 + \frac{qb(m-1)}{a+q(1-a)}) \\ &= \frac{1-b}{qb+1-b} + \frac{qbm}{qb+1-b} - (1 + \frac{qb(m-1)}{a+q(1-a)}) \\ &= 1 + \frac{qb(m-1)}{qb+1-b} - ((1 + \frac{qb(m-1)}{a+q(1-a)})) \\ &= qb(m-1)(\frac{1}{qb+1-b} - \frac{1}{a+q(1-a)}) \\ &= \frac{qb(m-1)}{(qb+1-b)(a+q(1-a))} (a+q(1-a) - (qb+1-b)) \\ &= \frac{qb(m-1)}{(qb+1-b)(a+q(1-a))} (q-1)(1-a-b) \\ &\leq 0 \end{aligned}$$

$$\begin{aligned} & u(\hat{\mathbf{p}}_5) - (1 + \frac{(q(1-a-b)+b)(m-1)}{b+q(1-b)}) \\ &= \frac{qa}{qa+(1-a)} + \frac{(1-a)m}{qa+(1-a)} - \left(1 + \frac{(q(1-a-b)+b)(m-1)}{b+q(1-b)}\right) \\ &= 1 + \frac{(1-a)(m-1)}{qa+(1-a)} - \left(1 + \frac{(q(1-a-b)+b)(m-1)}{b+q(1-b)}\right) \\ &= (m-1)(\frac{1-a}{qa+(1-a)} - \frac{q(1-a-b)+b}{b+q(1-b)}) \\ &= (m-1) \left(1 - \frac{qa}{qa+(1-a)} - (1 - \frac{qa}{b+q(1-b)})\right) \\ &= qa(m-1)(\frac{1}{b+q(1-b)} - \frac{1}{qa+(1-a)}) \\ &= \frac{qa(m-1)}{(qa+(1-a))(b+q(1-b))} (qa+(1-a) - b+q(1-b)) \\ &= \frac{qa(m-1)}{(qa+(1-a))(b+q(1-b))} (1-q)(1-a-b) \\ &\geq 0. \end{aligned}$$

So $A \leq u(\hat{\mathbf{p}}_6) - l(\hat{\mathbf{p}}_6)$, $B \leq u(\hat{\mathbf{p}}_5) - l(\hat{\mathbf{p}}_5)$, which implies

$$u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}}) \leq \max\{u(\hat{\mathbf{p}}_6) - l(\hat{\mathbf{p}}_6), u(\hat{\mathbf{p}}_5) - l(\hat{\mathbf{p}}_5)\},$$

i.e.,

$$R(f^{PAA}, \hat{\mathbf{p}}) \leq \max\{R(f^{PAA}, \hat{\mathbf{p}}_5), R(f^{PAA}, \hat{\mathbf{p}}_6)\}.$$

Putting the three pieces together, we have

$$R(f^{PAA}, \hat{\mathbf{p}}) \leq \max\{R(f^{PAA}, \hat{\mathbf{p}}_5), R(f^{PAA}, \hat{\mathbf{p}}_6)\},$$

which implies $\hat{\mathbf{p}}^*$ has exactly two non-zero entries.

Third Step: Next, we prove $\hat{\mathbf{p}}^* = (\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$. For $\hat{\mathbf{p}} = (a, 0, \dots, 0, 1-a), 0 \leq a \leq 1$, $l(\hat{\mathbf{p}}) = \frac{a}{a+q(1-a)} + \frac{q(1-a)}{a+q(1-a)} \times m$, $u(\hat{\mathbf{p}}) = \frac{qa}{qa+(1-a)} + \frac{1-a}{qa+(1-a)} \times m$, so

$$\begin{aligned} u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}}) &= (m-1) \times \left(\frac{1-a}{qa+(1-a)} - \frac{q(1-a)}{a+q(1-a)} \right) \\ &= (m-1) \times \frac{(1-q^2)(a-a^2)}{(1-q)^2(a-a^2)+q} \\ &= (m-1) \times \frac{(1+q)(a-a^2)}{(1-q)(a-a^2)+\frac{q}{1-q}}. \end{aligned}$$

Let $x = a - a^2 \in [0, \frac{1}{4}]$, $F(x) = \frac{u(\hat{\mathbf{p}}) - l(\hat{\mathbf{p}})}{m-1} = \frac{(1+q)x}{(1-q)x + \frac{q}{1-q}}$. Then

$$\begin{aligned} F'(x) &= \frac{(1+q) \left((1-q)x + \frac{q}{1-q} \right) - (1-q)(1+q)x}{\left((1-q)x + \frac{q}{1-q} \right)^2} \\ &= \frac{q(1+q)}{(1-q) \left((1-q)x + \frac{q}{1-q} \right)^2} \\ &> 0. \end{aligned}$$

So $x = \frac{1}{4}$ uniquely maximizes the regret. Since $x = \frac{1}{4} \iff a = \frac{1}{2}$, we prove $\hat{\mathbf{p}}^* = (\frac{1}{2}, 0, \dots, 0, \frac{1}{2})$.

Fourth Step: Finally, it is easy to figure out the optimal index for minimum $k_1(\hat{\mathbf{p}}^*)$ and the optimal index for maximum $k_2(\hat{\mathbf{p}}^*)$ can both be any value in $[m-1]$. Without loss of generality, we set $k_1(\hat{\mathbf{p}}^*) = 1, k_2(\hat{\mathbf{p}}^*) = m-1$. So the corresponding worst pair of information structure is

- $\theta_1 = (\mathbf{p}_1 = [\frac{1}{q+1}, 0, \dots, 0, \frac{q}{q+1}], \mathbf{g}_1 = [q, 1, \dots, 1, 1])$
- $\theta_2 = (\mathbf{p}_2 = [\frac{q}{q+1}, 0, \dots, 0, \frac{1}{q+1}], \mathbf{g}_2 = [1, 1, \dots, 1, q])$

□

THEOREM A.4. When $m = o\left(\left(\frac{n}{\ln n}\right)^{\frac{1}{4}}\right)$, PAA is $O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$ -optimal. That is, for any aggregator f ,

$$\max_{\mathbf{p}, \mathbf{g}} R(f^{PAA}, \mathbf{p}, \mathbf{g}) \leq \max_{\mathbf{p}, \mathbf{g}} R(f, \mathbf{p}, \mathbf{g}) + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right).$$

PROOF. For clarity, we define $R_n(f^{PAA}, \mathbf{p}, \mathbf{g})$, which is the regret of PAA when the sample size is n . Suppose f_n^* is the optimal aggregator when sample size is n . It is equivalent to prove

$$\max_{\mathbf{p}, \mathbf{g}} R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) \leq \max_{\mathbf{p}, \mathbf{g}} R_n(f_n^*, \mathbf{p}, \mathbf{g}) + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$$

Use $R_n(f^*)$ to denote the optimal regret when the sample size is n , and R^* is the optimal regret when $n \rightarrow \infty$. First Notice $R_n(f^*)$ is a non-increasing function with respect to n since we can always choose the aggregator which only uses part of the ratings, so $R_n(f^*) \geq R^*$. Then We bound the difference between $R_n(f^{PAA}, \mathbf{p}, \mathbf{g})$ and R^* by concentration inequality.

For any (\mathbf{p}, \mathbf{g}) , when $n \rightarrow \infty$, we obtain $\hat{\mathbf{p}} = \frac{\mathbf{p}\mathbf{g}}{\sum_j p_j g_j}$ and

$$R_\infty(f^{PAA}, \mathbf{p}, \mathbf{g}) = (f^{PAA}(\hat{\mathbf{p}}) - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}])^2 = \left(\frac{u(\hat{\mathbf{p}}) + l(\hat{\mathbf{p}})}{2} - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}]\right)^2$$

In the finite case, we get a noisy distribution $\hat{\mathbf{p}}'$. We use n_r to denote the number of people whose rating is r and \hat{n}_r to denote the number of people who actually report r .

$$R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) = \mathbb{E}[(f^{PAA}(\hat{\mathbf{p}}') - \frac{\sum_r r n_r}{n})^2] = \mathbb{E}[(\frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \frac{\sum_r r n_r}{n})^2]$$

By chernoff bound, we have

$$\Pr[|n_r - E[n_r]| \geq \sqrt{n \ln n}] \leq 2e^{-\frac{2n \ln n}{n}} = \frac{2}{n^2}$$

$$\Pr[|\hat{n}_r - E[\hat{n}_r]| \geq \sqrt{n \ln n}] \leq 2e^{-\frac{2n \ln n}{n}} = \frac{2}{n^2}$$

$$\Pr[|\sum_r \hat{n}_r - E[\sum_r \hat{n}_r]| \geq \sqrt{n \ln n}] \leq 2e^{-\frac{2n \ln n}{n}} = \frac{2}{n^2}$$

Define event $A_r = \{|n_r - E[n_r]| \geq \sqrt{n \ln n}\}$, event $A = \bigvee_{r=1}^m A_r$, event $B_r = \{|\hat{n}_r - E[\hat{n}_r]| \geq \sqrt{n \ln n}\}$, event $B = \bigvee_{r=1}^m B_r$, event $C = \{|\sum_r \hat{n}_r - E[\sum_r \hat{n}_r]| \geq \sqrt{n \ln n}\}$. By union bound, we have

$$\Pr[A \vee B \vee C] \leq \frac{4m+2}{n^2} \leq \frac{1}{n}$$

Case 1: A happens or B happens or C happens.

$$\mathbb{E}\left[\left(\frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \frac{\sum_r r n_r}{n}\right)^2 \mid A \vee B \vee C\right] = O(m^2)$$

Case 2: None of A, B, C happens. We have

$$\begin{aligned} |\hat{p}'_r - \hat{p}_r| &= \left| \frac{\hat{n}_r}{\sum_j \hat{n}_j} - \frac{p_r g_r}{\sum_j p_j g_j} \right| \\ &= \left| \frac{E[\hat{n}_r] + O(\sqrt{n \ln n})}{E[\sum_j \hat{n}_j] + O(\sqrt{n \ln n})} - \frac{p_r g_r}{\sum_j p_j g_j} \right| \\ &= \left| \frac{p_r g_r n + O(\sqrt{n \ln n})}{\sum_j p_j g_j n + O(\sqrt{n \ln n})} - \frac{p_r g_r}{\sum_j p_j g_j} \right| \\ &= O\left(\sqrt{\frac{\ln n}{n}}\right) \end{aligned}$$

When n is large enough, this error will not influence the optimal index. So

$$\begin{aligned} |l(\hat{\mathbf{p}}') - l(\hat{\mathbf{p}})| &= \left| \frac{\frac{1}{q} \sum_{j=1}^{k_1} j \hat{p}'_j + \sum_{j=k_1+1}^m j \hat{p}'_j}{\frac{1}{q} \sum_{j=1}^{k_1} \hat{p}'_j + \sum_{j=k_1+1}^m \hat{p}'_j} - \frac{\frac{1}{q} \sum_{j=1}^{k_1} j \hat{p}_j + \sum_{j=k_1+1}^m j \hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1} \hat{p}_j + \sum_{j=k_1+1}^m \hat{p}_j} \right| \\ &= \left| \frac{\frac{1}{q} \sum_{j=1}^{k_1} j \hat{p}_j + \sum_{j=k_1+1}^m j \hat{p}_j + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)}{\frac{1}{q} \sum_{j=1}^{k_1} \hat{p}_j + \sum_{j=k_1+1}^m \hat{p}_j + O\left(m \sqrt{\frac{\ln n}{n}}\right)} - \frac{\frac{1}{q} \sum_{j=1}^{k_1} j \hat{p}_j + \sum_{j=k_1+1}^m j \hat{p}_j}{\frac{1}{q} \sum_{j=1}^{k_1} \hat{p}_j + \sum_{j=k_1+1}^m \hat{p}_j} \right| \\ &= O\left(m^2 \sqrt{\frac{\ln n}{n}}\right) \end{aligned}$$

Similarly, $|u(\hat{\mathbf{p}}') - u(\hat{\mathbf{p}})| = O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$, so

$$\left| \frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \frac{u(\hat{\mathbf{p}}) + l(\hat{\mathbf{p}})}{2} \right| = O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$$

Notice

$$\left| \frac{\sum_r r n_r}{n} - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}] \right| = \left| \frac{\sum_r r n_r}{n} - \frac{\sum_r r \mathbb{E}[n_r]}{n} \right| = \left| \frac{\sum_r r (n_r - \mathbb{E}[n_r])}{n} \right| = O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)$$

Then we have

$$\begin{aligned} &\mathbb{E}\left[\left(\frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \sum_r r n_r\right)^2 \mid \neg(A \vee B \vee C)\right] \\ &= \left(\frac{u(\hat{\mathbf{p}}) + l(\hat{\mathbf{p}})}{2} - \mathbb{E}_{\mathbf{x} \sim \mathbf{p}}[\mathbf{x}] + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right)\right)^2 \\ &= R_\infty(f^{PAA}, \mathbf{p}, \mathbf{g}) + O\left(m^2 \sqrt{\frac{\ln n}{n}}\right) \end{aligned}$$

Putting the two pieces together, we have

$$\begin{aligned}
 R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) &= \Pr[A \vee B \vee C] \mathbb{E} \left[\left(\frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \frac{\sum_r r n_r}{n} \right)^2 \mid A \vee B \vee C \right] \\
 &+ \Pr[\neg(A \vee B \vee C)] \mathbb{E} \left[\left(\frac{u(\hat{\mathbf{p}}') + l(\hat{\mathbf{p}}')}{2} - \frac{\sum_r r n_r}{n} \right)^2 \mid \neg(A \vee B \vee C) \right] \\
 &= \Pr[A \vee B \vee C] O(m^2) + \Pr[\neg(A \vee B \vee C)] \left(R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right) \right) \\
 &\leq \frac{m^2}{n} + \left(1 - \frac{1}{n} \right) \left(R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right) \right) \\
 &= R_\infty(f^{PAA}, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right)
 \end{aligned}$$

Since for any \mathbf{p}, \mathbf{g} the inequality holds, we have

$$\max_{\mathbf{p}, \mathbf{g}} R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) \leq \max_{\mathbf{p}, \mathbf{g}} R_\infty(f^{PAA}, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right)$$

So for any sample size n ,

$$\begin{aligned}
 \max_{\mathbf{p}, \mathbf{g}} R_n(f^{PAA}, \mathbf{p}, \mathbf{g}) &\leq \max_{\mathbf{p}, \mathbf{g}} R_\infty(f^{PAA}, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right) \\
 &= R^* + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right) \\
 &\leq R_n(f^*) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right) \\
 &= \max_{\mathbf{p}, \mathbf{g}} R(f_n^*, \mathbf{p}, \mathbf{g}) + O \left(m^2 \sqrt{\frac{\ln n}{n}} \right)
 \end{aligned}$$

The first equality holds since f^{PAA} is the optimal aggregator when $n \rightarrow \infty$.

□

B Omitted Figures

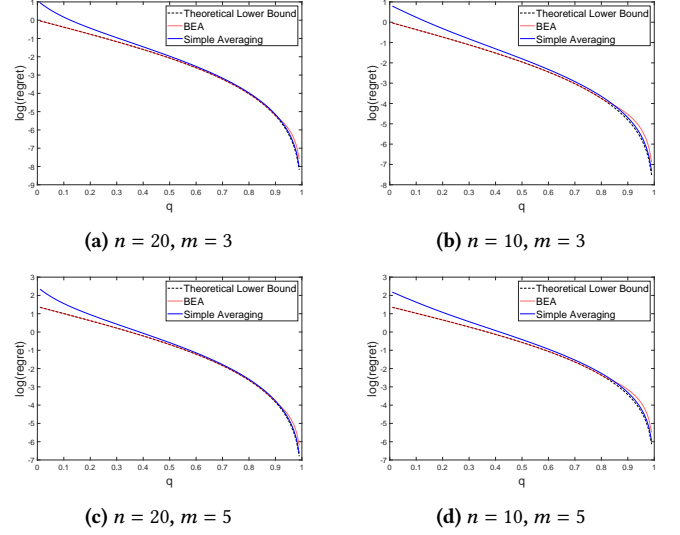


Figure 11: Simple averaging vs. BEA for different sample size n and the number of rating categories m . The x-axis is the lower bound of the participation probability, q , and the y-axis is the natural logarithm of the regret. The regret of BEA almost matches the theoretical lower bound for a wide range of q .

Figure 11 are the regret of different aggregators with q ranging from 0.01 to 0.99 with step 0.01. Notice when q is large (the threshold is around 0.82 for $n = 10$ and 0.91 for $n = 20$), the regret of BEA deviates from the theory lower bound and even exceeds the regret of simple average aggregator when q is larger. This is because the worst information structure for BEA changes from the pair in Lemma 3.2 to

$$\theta = (\mathbf{p} = [b, 0, \dots, 0, 1 - b], \mathbf{g}_1 = [q, q, \dots, q, q])$$

where b is dependent with n and q .

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009