

Incentivizing Effort and Precision in Peer-grading

Anujit Chakraborty¹, Jatin Jindal², and Swaprava Nath³

¹University of California Davis, chakraborty@ucdavis.edu

²Zamstars, jatinjindal369@gmail.com

³Indian Institute of Technology Kanpur, swaprava@iitk.ac.in

Abstract

Most peer-evaluation practices rely on the peer evaluator’s goodwill for an accurate evaluation. But what if graders are competitive, i.e., enjoy higher utility when their peers get lower scores? We introduce a new mechanism, PEQA, that uses a *score-assignment rule* and a *grading performance score* to incentivize such graders. The score-assignment rule aggregates multiple peer-evaluations to assign a final score that is free from grader-bias. The rule and the performance score also guarantee that any grader’s utility increases monotonically with her grading precision, irrespective of her competitiveness and how her co-graders act. In a class of assignment rules that generalize over the least squares estimator, our assignment rule uniquely satisfies this utility-precision monotonicity while allowing flexibility in how large performance scores need to be. When grading costs are private information, a special version of PEQA implements the socially optimal effort-choices in an equilibrium of the peer-evaluation game among the co-graders. Data from our classroom experiments confirm our theoretical assumptions and show that PEQA outperforms the median mechanism, which is popular in Massive Open Online Courses.

Keywords: peer evaluation, massive open online courses, mechanism design, bias insensitivity, reliability monotonicity, experiment.

1 Introduction

A peer-evaluation process aggregates assessments from peers to judge the quality of submitted work. Scientific communities use peer-evaluation for reviewing the quality of articles and grant proposals (Campanario, 1998). Massive Open Online Courses (MOOCs) and in-person classes have recently started adopting peer-grading too, for three reasons. First, it simplifies and accelerates the evaluation and grading process. Second, it improves learning outcomes of participating students (Sadler and Good, 2006). Third, it easily scales to large classes.

There is, however, a scope for skepticism about the accuracy of peer-graded outcomes. A grader might be unmotivated to evaluate diligently when peer-grading is effort-intensive and unincentivized. She might also be biased in her evaluations if she cares about her relative success within peers. Strong et al. (2004) find that peer-graders give consistently biased grades in peer-grading schemes. In an anonymous survey that we ran on the students of the Indian Institute of Technology Kanpur, 49% of the 549 respondents expected that their fellow students would grade aggressively to reduce the scores of others, and thereby try to improve their relative class-ranking.

We model peer-grading of exams as a game with strategic and competitive subjects. Subjects (graders) write an exam and then peer-grade each others’ exams. Their utility is linearly increasing in their total course-scores and linearly decreasing in their peer’s total course-scores. The total course-score is the sum of their own exam score (aggregated from peer-reports) and a score based on their peer-grading performance.

To model the grading and reporting process, we generalize the \mathbf{PG}_1 statistical peer-grading model of Piech et al. (2013)¹. We drop its parametric distribution assumptions for the sake of generality and model grading as a strategic decision. Our peer-graders observe an independent, noisy signal about the true score of each paper they are asked to grade. The signal comes from any differentiable distribution whose mean is at the true score and support is $(-\infty, \infty)$. The graders can choose the precision (variance) of the signal-distribution: higher effort produces more precise signals. Graders can then add a bias while reporting their observed signal. Strategic graders maximize their utility by choosing their effort-level (reflected in the signal-variance) and bias.

We propose a new peer-grading mechanism, Peer Evaluation with Quality Assurance (PEQA), that determines how graders are assigned to papers, how final scores are evaluated from peer-reports, how grading performance scores are handed out, and provides option to the students to raise regrading requests. Each grader is assigned to $K > 2$ papers and they never grade their own paper. A subset of the total number of papers (called *probes*) is evaluated by the teaching staff. The peer-graders cannot tell apart the probes from the non-probes. PEQA compares the grader’s and the teaching staff’s evaluations of the probes to estimate each grader’s bias and precision. This requires two identifying assumptions: that the teaching staff can observe the true scores on the probe papers, and, that the graders grade identically on probes and non-probes. The estimated grader-bias is subtracted from the peer-reports to de-bias the reports. PEQA’s score-assignment function assigns a weighted average of the de-biased grader-reports, with the weights being the inverse square-root of the estimated grader-variance. Thus, reports from high variance graders play a smaller role in the finally assigned score.

We allow students to raise regrading requests after seeing their score. The teaching staff regrade such papers and assign them the true score. We assume that such requests are raised only when the student knows that her initially assigned score was lower than the true score.

PEQA derives grading performance scores from its score-assignment function, using a transfer that resembles the Vickrey-Clarke-Groves (VCG) transfer (Vickrey, 1961; Clarke, 1971; Groves, 1973). We calculate the assigned score of paper j once for its whole group of peer-graders $G(j)$, and then for $G(j) \setminus \{i\}$, i.e., the group after removing grader i . We calculate the expected accuracy of the assigned scores with respect to the true scores, in both scenarios. Grader i is awarded a performance-score that is equal to the difference between the accuracies of groups $G(j)$ and $G(j) \setminus \{i\}$. Thus i ’s performance scores can be interpreted as her marginal contribution to the expected accuracy in score-assignment (details in Algorithm 1). Linking grading-incentives to marginal contribution motivates accuracy and collective efficiency, as we discuss later. How do we calculate the accuracy² of assigned scores? Squared distance from the true score is one potential measure of accuracy, but our grading performance scores work equally well for a more general class of measures.

We show that graders have no incentives to bias their reports under PEQA. This is because the score-assignment function de-biases the grades through the estimated grader bias. Thus, i ’s bias cannot lower others’ assigned final scores (eliminating any competitive incentive). Further, i ’s grading performance score depends on the debiased score-assignment function, and hence is unaffected by bias. Though the bias estimates from probes are noisy, they are correct in expectation.

Under PEQA, the utility of any grader is strictly increasing in her precision (and thus effort), despite her competitive preferences and irrespective of how the co-graders act. Further, this holds for any realizations of true scores on papers and the noise in other’s reports (Theorem 1). Thus, the noisier the grading report by grader i , the lower is her expected utility. Two channels are driving this result. First, for the group $G(j)$ who grade paper j , a lower grading precision of

¹This is a widely used model for peer evaluation in MOOCs.

²We have used two relatively similar terms *precision* and *accuracy*. A grader chooses the *precision* of her signal, while *accuracy* reflects how close the mechanism-assigned score is to the true score.

$i \in G(j)$ lowers her marginal contribution to accurate score-assignment on paper j . This lowers i 's performance score since it is equal to her marginal contribution. Second, grader i 's noisier grading leads to a noisier assigned grade on paper j . The noise moves the assigned grade on paper j above or below its true grade. Higher is the noise, larger is the potential movement in either direction. Grader i can determine the magnitude of the noise through her effort, but not the direction in which the noise moves the assigned grade on paper j . By selectively raising regrading requests, student j keeps any undeserved high grades and removes any undeserved low grades caused by i 's noisy grading. Thus, i 's imprecise grading ends up increasing j 's grades, post regrading-requests. Given i dislikes when j gets higher grades, this decreases i 's utility in expectation. Thus, i 's competitiveness also fuels i 's desire for precise grading!

The score-assignment function under PEQA aggregates information efficiently in a world with incentive concerns. Without incentive or strategic concerns, a statistician would suggest a score-assignment function that that would minimize the expected squared distance between the assigned score and true score on exam j , conditional on the true bias and variance parameters. Then, those true parameters could be approximated by the estimated bias and variance. In Equation (26) of Appendix B, we show that under the strong distribution-assumptions of Piech et al. (2013), such a score-assignment function on exam j would look like a weighted-aggregator:

$$r_j = \frac{\lambda_0 \times \text{prior mean of all exams} + \sum_i \lambda_i \times \text{de-biased report on exam } j \text{ from grader } i}{\lambda_0 + \sum_i \lambda_i} \quad (1)$$

with the weights λ_i equal to the inverse of grader-variance. PEQA uses a similar assignment function but with weights equal to the inverse *square-root* of grader-variance. This difference is explained by our two major objectives: (1) modeling grading as a strategic choice, and (2) incentivizing accurate grading under flexible performance scores. If the course-designer wants to use low performance scores in her course, the competitiveness of graders must also be utilized to fuel the incentives for precision. Our weighted aggregation rule uniquely guarantees (Theorem 2) that competitiveness fuels precision, irrespective of other's actions and noise terms, and irrespective of distribution assumptions. Statistically motivated aggregation rules are unable to use competitiveness to incentivize precision.

By linking incentives from grading to marginal contribution, PEQA incentivizes socially efficient effort-allocation. When grading is costly, increasing accuracy of the mechanism-assigned exam score comes at an increased individual (and social) cost of effort. We formulate peer-grading as a game, where each subject independently decides how much effort to put in, subject to her personal costs and benefits. We define social value as the total grading accuracy minus total effort cost, aggregated over all papers and graders. When the cost of grading is a private information, a centralized assignment of socially optimal effort-allocation is impossible. Further, each grader's effort influences the marginal effectiveness (and hence performance score) of her co-graders, creating an externality. We propose a special case of the grading performance score function, that pays each grader her marginal social contribution and compensates her for disutility from competitiveness, without altering the relative ranking of subjects by their total score. Under this special version of the PEQA mechanism, there are Nash equilibria of the peer-grading game in which individuals spend effort such that the social value is maximized (Theorem 3).

We ran classroom experiments (Section 7) to test some of our baseline assumptions and to see how easily our mechanism could be implemented in practice. Students enrolled in a computing course were asked to peer-grade a weekly class-quiz. To induce a competitive environment, we paid subjects by the relative ranking of their total scores in class. Thus, payments would weakly increase and decrease in own and other's total scores respectively. The total scores were the sum of their peer-evaluated score and grading performance score, as in the model. Subjects were explained how PEQA assigns final scores and grading-performance scores. We graded all

the quizzes ourselves to compare evaluations under PEQA to our own. The scores assigned under PEQA were remarkably precise, and only 1 out of 41 had a wrong score.

Our PEQA sessions confirm two of our assumptions. First, the bias and variance were indeed identical across probes and non-probes: subjects were not able to discern one from the other. Second, grader-manipulations, whenever present, reduced scores instead of inflating scores. This rejects the existence of collusive (i.e., the opposite of competitive) graders.

We ran a second competitive session under a Median mechanism, which is currently the most popular mechanism used in MOOCs. Subjects were paid similarly to PEQA sessions, i.e., according to their relative ranking of the total scores, but there was no grading performance score added to the total score. We used the median of peer reports as final and total scores. We equated the payments made to subjects in PEQA and Median sessions for each quartile, so that both mechanisms paid identically. We used the same class-quiz as before. Students had to grade the same number of papers as in the PEQA sessions. In our experiments, PEQA mechanism outperformed Median mechanism in terms of allocating precise final scores. These differences were statistically significant.

1.1 Related Work

The existing research on peer-evaluation mechanisms can be broadly divided into three strands. The first strand of literature abstracts away from any strategic motives of the peer-evaluators. Instead of providing a mechanism to incentivize strategic evaluators, they propose how the grader reports could be aggregated efficiently (Hamer et al., 2005; Cho and Schunn, 2007; Piech et al., 2013; Shah et al., 2013; Paré and Joordens, 2008; Kulkarni et al., 2014; De Alfaro and Shavlovsky, 2014; Raman and Joachims, 2014; Caragiannis et al., 2015; Wright et al., 2015).

The second strand of literature is based on *peer-prediction* approaches. These mechanisms incentivize *coordination* on similar evaluation reports by punishing evaluations that don't match each other. Thus, they do not necessarily incentivize *accuracy* (Prelec, 2004; Miller et al., 2005; Jurca and Faltings, 2009; Faltings et al., 2012; Witkowski et al., 2013; Dasgupta and Ghosh, 2013; Witkowski and Parkes, 2013; Waggoner and Chen, 2014; Shnayder et al., 2016). Any such mechanism introduces uninformative equilibria alongside the truth-telling one (Jurca and Faltings, 2009; Waggoner and Chen, 2014).³ More recent developments make the truthful equilibrium Pareto dominant, i.e., the truthful equilibrium is (weakly) more rewarding to every agent than any other equilibrium (Dasgupta and Ghosh, 2013; Witkowski and Parkes, 2013; Kamble et al., 2015; Radanovic and Faltings, 2015; Shnayder et al., 2016).

The final strand consists of *hybrid* approaches where the true quality of some of the peer-assessed material can be found, for e.g. via evaluating a part of the materials by the mechanism designer (teaching staff in case of MOOCs) herself. Graders are then rewarded for agreement with the designer-agreed report (Jurca and Faltings, 2005; Dasgupta and Ghosh, 2013; Gao et al., 2016). Our mechanism also utilizes the feature that the true scores on a small subset of assignments can be revealed at a small cost. However, additionally, we address new and practical features of the peer-grading problem: we allow for competitive graders, we solve the efficient allocation problem under costly grading, and we allow regrading requests.

Alon et al. (2011) and Holzman and Moulin (2013) study situations where peers have to choose a subset amongst themselves for a reward. The challenge here is to incentivize the peers to reveal their private information unselfishly. In particular, the goal is to guarantee that what peers report does not affect their chances of winning or getting selected. In these settings, there is no need to incentivize peers to gather information that is 'objective' (e.g., true score on an exam) and verifiable at a cost. There is also no need to ensure that peers enjoy higher utility when their gathered information is more precise. Finally, peers are purely selfish: they do not

³In particular, when the information is costly to obtain, it is generally easier for the agents to resort to coordinating on an uninformative low-effort equilibrium.

care about who wins in case they do not win themselves. Thus, by debriding the reports from personal winning chances, the mechanism makes the peers indifferent between all reports.

Cai et al. (2015) consider a setting where data-sources (e.g., human labelers) can be paid monetarily to get their estimation of $f(x_i)$ at points x_i allocated to them. The end goal is to estimate an exogenously provided f using a given estimator \hat{f} . Data-sources can observe a noisy version of $f(x_i)$ with the noise decreasing in their effort, and they maximize the difference between the payment and the cost of the effort. They show that under their VCG-like payment mechanism and the assumption of a “well-behaved” \hat{f} , the dominant strategy for a data-source is to reveal its observation correctly and always participate in the data-providing exercise. Cai et al. (2015)’s data-sources naturally have no competitive preferences, like our graders do. We also propose the optimal estimator \hat{f} to use on the observed data, which they don’t.

2 Setup of peer-grading Mechanism

Each subject $i \in N = \{1, \dots, n\}$ has written an exam, and is also a participant in the peer-grading process. Thus $N = \{1, \dots, n\}$ represents both the set of answer-scripts to be graded and the set of graders. We use i as the index for a grader and j as the index for a paper.

A fixed number $\ell (<< n)$ of these papers are evaluated by the teaching staff, and the true grades on them are thus known. These papers are called the *probe* papers. Let $G(j)$ denote the set of peer-evaluators of paper j and $G^{-1}(i) := \{k \in N : i \in G(k)\}$ denote the set of papers assigned to evaluator i . The set $P_i \subset G^{-1}(i)$ and $NP_i = G^{-1}(i) \setminus P_i$ denotes respectively the probe and non-probe papers assigned to i . Both true and reported scores belong to the set of scores S which is an interval in \mathbb{R} . The co-graders of individual i are $CG_i = \cup_{j \in NP_i} G(j) \setminus \{i\}$. We assume that the co-graders of i grade at least one common non-probe paper with i .

A *peer-grading mechanism* M is the tuple $\langle G, \mathbf{r}, \mathbf{t} \rangle$, where

- ▷ G is the *assignment* function $G : N \rightarrow 2^N$ that maps papers to graders.
- ▷ $\mathbf{r} : \times_{j \in N} S^{G(j)} \rightarrow S^n$ is the *score-assignment* function, where the j th component $r_j(\cdot)$ is the function assigning the final score of paper j based on the scores reported by $G(j)$.
- ▷ $\mathbf{t} : \times_{i \in N} S^{G^{-1}(i)} \rightarrow S^n$ is the *peer-grading performance score* function, where the i th component $t_i(\cdot)$ is the function that yields the peer-grading performance score to grader i .

2.1 Model of the True and Reported Scores

We generalize the following \mathbf{PG}_1 model of true score, bias, and reliability (Piech et al. (2013)). We make two major changes. First, we replace their assumptions of normality with a general distribution $\mathcal{F}(\cdot)$ with a support of $(-\infty, \infty)$ and differentiable density function $f(\cdot)$. We use $\mathcal{F}(\mu, 1/\gamma)$ for such a distribution with mean μ and variance $1/\gamma$. Second, instead of assuming that bias and reliability are drawn randomly and independently from a Normal and Gamma distribution, we make each a strategic choice by the peer-graders. Subject to these changes, the following features in our model resemble the \mathbf{PG}_1 model.

- ▷ The true score y_j for paper j is distributed as $\mathcal{F}(\mu, 1/\gamma)$, for all $j \in N$. This distribution is known from historical data of past examinations.
- ▷ $\tilde{y}_j^{(i)}$ is the reported score of paper j by grader i . It is distributed as $\mathcal{F}(y_j + b_i, 1/\tau_i)$, where b_i and τ_i are called the *bias* and *reliability* of i respectively.
- ▷ We have used the same distribution \mathcal{F} for both the true scores y_j and the score observed by the grader i , i.e., $\tilde{y}_j^{(i)}$, to keep the model similar to \mathbf{PG}_1 . However, this is not critical to our results. In particular, (a) we can have two different distributions for these two sets of random variables, and (b) the distribution of the observed score $\tilde{y}_j^{(i)}$ can vary with i . None of these will affect the main conclusions of this paper.

Bias is the grader’s *systematic deviation from the true score* and reliability is her *effort-driven precision*. Reliability could be interpreted as the inverse of noise. Bias could originate from strategic manipulation or from non-strategic (generous or strict) grading-habits. Thus a grader might or might not get to choose bias, but she chooses her reliability. In this paper, we would assume that the grader chooses her bias. Later (in Theorem 1) we would show that our mechanism is robust to grading bias, irrespective of its strategic or non-strategic nature.

We assume that a grader grades all papers (probes and non-probes) with the same bias and reliability. This assumption is natural if the graders cannot tell the probes from the non-probes. We find support for this in our experimental sessions. Bias and reliability are indeed identical across probes and non-probes.⁴ We use the shorthand $\theta_i = (b_i, \tau_i) \in \mathbb{R} \times \mathbb{R}_{\geq 0}$ to denote grader i ’s strategic choices.

2.2 Other primitives of our mechanism

In this section, we define $\langle G, \mathbf{r}, \mathbf{t} \rangle$ functions to describe our mechanism.

Paper assignment rule $G^*(\cdot)$

Every paper is graded by at least one grader, and every grader grades at least two probe and one non-probe papers. Thus, (a) $G^*(j) \neq \emptyset$ and $j \notin G^*(j)$, $\forall j \in N$, (b) $|P_i| \geq 2$, $\forall i \in N$, and (c) $NP_i \neq \emptyset$, $\forall i \in N$. The graders know the proportion of probe and non-probe papers assigned to them, but cannot tell them apart.

Grade assignment and performance scores

The mechanism compares the peer-graded scores ($\tilde{y}_j^{(i)}$) with true scores (y_j) on the probe papers P_i , to statistically estimate the error parameters $\hat{\theta}_i = (\hat{b}_i, \hat{\tau}_i) \in \mathbb{R} \times \mathbb{R}_{\geq 0}$ of each grader i . We have $\hat{b}_i = \frac{\sum_{j \in P_i} (\tilde{y}_j^{(i)} - y_j)}{|P_i|}$ and $\hat{\tau}_i = \frac{|P_i| - 1}{\sum_{j \in P_i} (\tilde{y}_j^{(i)} - (y_j + \hat{b}_i))^2}$. The estimated parameters are used in assigning performance-scores to papers and performance scores to peer-graders.

DEFINITION 1 (Score and Reward) *We define the score-assignment rule and the social reward as follows.*

▷ *The score-assignment function $\mathbf{r} = (r_j : j \in N)$ is inverse standard-deviation weighted de-biased mean (ISWDM) if for every non-probe paper j , it assigns*

$$r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}) = \frac{\sqrt{\gamma}\mu + \sum_{i \in G(j)} \sqrt{\hat{\tau}_i} (\tilde{y}_j^{(i)} - \hat{b}_i)}{\sqrt{\gamma} + \sum_{i \in G(j)} \sqrt{\hat{\tau}_i}}, \quad (2)$$

where $\tilde{y}_j^{(i)}$ is the evaluation by the i th peer-grader and $(\hat{b}_i, \hat{\tau}_i)$ are her estimated parameters. Score \mathbf{r}^* assigns the instructor-verified grade on every probe paper.

▷ *The social reward for paper j , at a score r_j^* and true score y_j , is*

$$W_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}, y_j) = R(r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j), \quad (3)$$

where $\tilde{\mathbf{y}}_j^{G(j)}$ is the vector of peer-evaluated scores reported on paper j , $\hat{\boldsymbol{\theta}}_{G(j)}$ is the vector of evaluated error-parameters for the relevant graders $G(j)$, and $R(\cdot)$ is the reward function defined as follows.

⁴We suggest measuring bias and reliability question-wise. This assumes that a grader has a fixed bias and reliability on the same question, for all the papers she grades.

We assume that $R(x, x) = 0 \geq R(x, y) = R(y, x)$ for all $x, y \in S$ and $R(x_1, y_1) < R(x_2, y_2)$ if $|x_1 - y_1| > |x_2 - y_2|$ for all $x_1, x_2, y_1, y_2 \in S$. One example of such a function would be $R(x, y) = -(x - y)^2$, which calculates the squared error in assigned scores.

▷ The social reward at a score r_j^* for paper j without grader i when the true score is y_j is denoted by $W_j^{(-i)*} = W_j^*(\tilde{\mathbf{y}}_j^{G(j)\setminus\{i\}}, \hat{\boldsymbol{\theta}}_{G(j)\setminus\{i\}}, y_j)$ where $W_j^*(\cdot)$ is defined as above.

The parameters γ, μ, b_i , and τ_i are as given by the **PG**₁ model of Piech et al. (2013) (see Section 2.1).

We will use the shorthands W_j^* and $W_j^{(-i)*}$ for the above two expressions when the arguments of such functions are clear from the context. The ISWDM score-assignment function takes a *weighted average* of the prior mean μ and the de-biased (subtracting the estimated bias from the reported scores) reported scores. De-biasing ensures that the biases of the graders do not affect the finally assigned grade. The weight is chosen to be the square-root of reliability, which is the inverse of the variance for that grader. Higher the estimated precision, higher is the weight on a grader.

One could imagine a more general class of *weighted average (WA) score-assignment functions*

$$r_j^{\text{WA}}(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}) = \frac{\lambda_0 \mu + \sum_{i \in G(j)} \lambda_i (\tilde{y}_j^{(i)} - \hat{b}_i)}{\lambda_0 + \sum_{i \in G(j)} \lambda_i}, \quad (4)$$

where $\lambda_0, \lambda_i \geq 0, \forall i \in N$, not all zero. The weight λ_i is assumed to be *independent* of σ_{-i} , $\forall i \in N$, so that, the relative weight of any two graders is independent of the reliabilities of the other graders.

One could instead have used this score-assignment function and the reward could be computed accordingly. In particular, a more natural score-assignment function is one that minimizes the mean squared error in assigned scores. When the social reward is defined as the negative of the mean squared error, this is indeed the *expected (social) reward maximizer (ERM)*:

$$r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}) \in \operatorname{argmax}_{x_j \in S} \mathbb{E}_{y_j | \tilde{\mathbf{y}}_j^{G(j)}; \hat{\boldsymbol{\theta}}_{G(j)}} R(x_j, y_j). \quad (5)$$

The expression of r_j^{ERM} for the Piech et al. (2013) grader error model (Section 2.1) with quadratic reward function reduces to an expression same as Equation (4) with $\lambda_0 = \gamma$ and $\lambda_i = \hat{\tau}_i, \forall i \in N$ (see Appendix B for details). However, in Theorem 2 we will show that in the class of score-assignment functions, ISWDM uniquely satisfies some desirable properties.

Regrading Requests

We consider peer-grading mechanisms that allow regrading requests. We assume that when a regrading request is raised, the instructor regrades the paper herself and assigns the true score on the paper. We also assume that the students know the true scores on their own papers and only raise a regrading request when they expect it to raise their score further.

ASSUMPTION 1 Student j knows y_j and raises a regrading request only if $r_j^* \leq y_j$.

3 The PEQA mechanism

Algorithm 1 shows the detailed steps of PEQA. It assigns probe and non-probe papers to the graders. The assigned score on a paper is a weighted average (with appropriately chosen weights). Finally, the grading performance score is the *marginal contribution* of the grader towards the social reward. In the next section, we discuss some desirable properties of a peer-grading mechanism.

Algorithm 1 PEQA

- 1: Inputs: (1) the parameters μ and γ of the priors on y_j , $\forall j \in N$, which is distributed as $\mathcal{F}(\mu, 1/\gamma)$, (2) the reported scores $\tilde{\mathbf{y}}_P^N$ of the graders on the probe papers, and (3) reported scores $\tilde{\mathbf{y}}_{N \setminus P}^N$ on the non-probe papers.
 - 2: Set the probe set P with $|P| = \ell$, a pre-determined constant $\leq \frac{n}{\frac{K}{2} + 1}$, where K (even) is the number of papers assigned to each grader.
 - 3: $G = G^*$: every grader $i \in N$ is assigned $\frac{K}{2}$ probe and $\frac{K}{2}$ non-probe papers, in such a way that every non-probe paper is assigned to at least $\frac{K}{2}$ and at most $\frac{K}{2} + 1$ graders. This is always possible by assigning the $(n - \ell)$ non-probe papers to $(n - \ell)$ graders with each paper assigned to exactly $\frac{K}{2}$ graders. The rest ℓ graders can be assigned to the same $(n - \ell)$ papers arbitrarily such that these papers get at most one additional grader (since $\ell \frac{K}{2} \leq n - \ell$). Note that this is the reason ℓ cannot be larger than $n/(K/2 + 1)$. Ensure that a grader does not get her own paper for evaluation.
 - 4: Estimate $\hat{b}_i, \hat{\tau}_i, \forall i \in N$ as given in Section 2.2.
 - 5: \mathbf{r} : the score of the paper j is given by the ISWDM \mathbf{r}^* (Equation (2)).
 - 6: At this stage, students may request for regrading. Instructor learns the correct grade y_j for the papers which came for regrading. For other papers, $y_j = r_j^*$ is assumed.
 - 7: \mathbf{t} : the performance score to grader i for grading paper $j \in NP_i$ is given by $t_i^j = \alpha(W_j^* - W_j^{(-i)*})$, where $\alpha > 0$ is a constant chosen at the designer's discretion. The total performance score to grader i is therefore $t_i = \sum_{j \in NP_i} t_i^j$.
-

4 Incentives and Design Desiderata

Individual Preferences We assume that every individual i cares about (1) her total score (sum of her exam score and peer-grading performance score t_i), and (2) the total scores of other individuals. Her utility is linear, and it is increasing in (1), weakly decreasing in (2). For agent i in mechanism $M = \langle G, \mathbf{r}, \mathbf{t} \rangle$, the utility is given by

$$u_i^M = r_i + t_i - \left(\sum_{j \in N \setminus \{i\}} w_{ij} \cdot (r_j + t_j) \right), \quad (6)$$

where $w_{ij} \geq 0$.

Individuals are aware of the score-assignment rule, social reward (Definition 1), and regrading procedure (Assumption 1) under Algorithm 1. They also understand how regrading rules and Algorithm 1 interact with their grading incentives. For example, they know that if based on their reports, the mechanism assigns a generous score $r_j^* \geq y_j$, then no regrading requests would be raised and the assigned score would be interpreted as the true score. Further, the social reward for their group would be calculated to be at the highest possible value. Further, PEQA would calculate the counterfactual social reward $W_j^{(-i)*}$ and their grading performance scores t assuming that the assigned score is the true score. In all other cases, a regrading request would be filed, the actual true score would be investigated by the teaching team, and then the social rewards would be calculated with respect to it.

A few uncertainties are resolved after the grader i (and her co-graders) chooses her decision variables (b_i, τ_i) and before r^* and t^* are computed. These are, (1) the scores finally reported by this grader i , which on any paper j , come from $(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b_i, 1/\tau_i)$, (2) the decision variables (b_k, τ_k) chosen by co-grader k (i.e., the strategic uncertainty), (3) the true score y_j on paper j that determines if regrading requests are made and what the incentives from grading are, (4) the score on paper j reported by a co-grader k , which follow the distribution $(\tilde{y}_j^{(k)} | y_j) \sim$

$\mathcal{F}(y_j + b_k, 1/\tau_k)$.

We define two desirable properties of peer-grading mechanisms. The properties consider the grader i 's expected utility from the choice of strategies she makes. *All expectations are only being taken with respect to uncertainty (1), the distribution of i 's grade-evaluation process $(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b_i, 1/\tau_i)$.* The properties hold for any ex-post realization of the other uncertainties (2) to (4), and there is no expectation taken on them. This is why both properties are defined as *ex-post*.

DEFINITION 2 (Ex-Post Bias Insensitivity (EPBI)) *A peer-grading mechanism $M = \langle G, \mathbf{r}, \mathbf{t} \rangle$ is ex-post bias insensitive, if the expected score assigned to any paper j and the expected utility of participant i , are independent of the bias b_i . Bias independence holds irrespective of the biases and reliabilities of other graders $j \neq i$, the true score y_j , and reported scores of the other graders. Mathematically,*

$$E_{(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b_i, 1/\tau_i)} u_i^M(\tilde{y}_j^{(i)}, \tilde{\mathbf{y}}_j^{(-i)}, y_j) = E_{(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b'_i, 1/\tau_i)} u_i^M(\tilde{y}_j^{(i)}, \tilde{\mathbf{y}}_j^{(-i)}, y_j),$$

$$\forall \{\tilde{y}_j^{(k)}, b_k, \tau_k\}_{k \neq i}, \forall y_j, \tau_i, \forall j \in G^{-1}(i), \forall i \in N. \quad (7)$$

DEFINITION 3 (Ex-Post Reliability Monotonicity (EPRM)) *A peer-grading mechanism $M = \langle G, \mathbf{r}, \mathbf{t} \rangle$ is ex-post reliability monotone if for every grader, the utility is monotonically increasing with her reliability, irrespective of the biases and reliabilities chosen by the other graders, the realizations of the true scores and the scores reported by the different graders. Mathematically,*

$$E_{(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b_i, 1/\tau_i)} u_i^M(\tilde{y}_j^{(i)}, \tilde{\mathbf{y}}_j^{(-i)}, y_j) > E_{(\tilde{y}_j^{(i)} | y_j) \sim \mathcal{F}(y_j + b'_i, 1/\tau'_i)} u_i^M(\tilde{y}_j^{(i)}, \tilde{\mathbf{y}}_j^{(-i)}, y_j),$$

$$\forall \tau_i > \tau'_i, \forall \{\tilde{y}_j^{(k)}, b_k, \tau_k\}_{k \neq i}, \forall y_j, b_i, \forall j \in G^{-1}(i), \forall i \in N. \quad (8)$$

Both these properties are stronger than *dominant strategies*. The utility depends on the realizations of random variables like the true scores \mathbf{y} and the reported scores $\tilde{\mathbf{y}}$. A dominant strategy definition will only require the (in)equalities to be satisfied after taking *interim* expectation over some relevant distribution over those variables. However, our *ex-post* properties require this to be satisfied for *every* realization of these random variables.

5 Properties of PEQA

All results in this section hold under a general model of how peer-reports are generated. For example, we assume that signals come from a general distribution $\mathcal{F}(\cdot)$ that has a *support of* $(-\infty, \infty)$.

Our first result shows that PEQA satisfies both the properties mentioned in Section 4, as long as the subjects care more than their own scores than about others' scores. The result holds for any grading performance score, as long as it pays according to marginal contributions to some social reward function (see Definition 1 and Algorithm 1). We defer the proofs of the results to Appendix A.

THEOREM 1 *If $\sum_{k \in N \setminus \{i\}} w_{ik} \leq 1, \forall i \in N$, then PEQA is EPBI and EPRM.*

A direct consequence of this result is that a grader will have no incentive in putting a deliberate upward or downward bias in this competitive environment and also will find it in her interest to maximize her precision (subject to her effort-costs).

The EPBI result is driven by how the score-assignment function de-biases the grades through the estimated grader bias. Though the bias estimates from probes are noisy, in expectation, they are correct and are identical across probes and non-probes. Thus grader i 's bias cannot lower other's assigned final scores. We show that bias also does not effect the post-regrading expected score $\max\{r_j^*(\cdot), y_j\}$. Thus biasing reports does not provide any competitive incentives. Her grading performance score depends only on the assigned final scores on the papers she graded, and hence it is unaffected by bias too. EPBI is independent of the condition on w_{ik} s.

There are two channels driving the EPRM result.

- ▷ The first channel derives from how we define our score-assignment function *and* how our grading performance scores are equal to marginal contributions. Under our score-assignment function, grader i 's noisier grading leads to a noisier assigned grade on paper j .⁵ Thus, a lower grading precision of $i \in G(j)$ invariably lowers i 's marginal contribution to accurate score-assignment on paper j . This lowers i 's grading performance score, which is equal to marginal contribution under our mechanism.
- ▷ The second channel depends only on the score-assignment function and our regrading assumption (see Assumption 1). As mentioned previously, under our score-assignment function, grader i 's noisier grading leads to a noisier assigned grade on paper j . The noise moves the assigned grade above or below the true grade. Higher is the noise, larger is the potential movement in either direction. Grader i determines the magnitude of the noise, but not the direction in which the noise moves the assigned grade. By selectively asking for regrades, student j keeps any undeserved high grades and reverses any low grades that result from the noise. Thus, i 's noisier grading ends up increasing j 's grades post regrading-requests. Given i dislikes when j gets higher grades, this decreases i 's utility in expectation. Thus, i 's competitiveness also fuels i 's desire for an accurate grading.

Deriving the EPRM condition requires a bound on w_{ik} s. This is because the choice of precision of grader i affects the final grades of whoever she grades, and the marginal contributions (thus, grading performance scores) of her co-graders. It is the competitiveness with respect to the latter that requires the w_{ik} -condition. In particular, we show that the condition on w_{ik} s is sufficient to ensure that the collective weight on other's grading performance scores never outweighs a competitive student's regard for her own performance score, irrespective of other's actions and noise.

The w_{ik} condition can be further weakened to the following condition: The sum of the weights ($\sum_{k \in A} w_{ik}$) on any $K/2$ potential co-graders does not outweigh i 's concern about her own performance score score in her utility.⁶

The weight α that an instructor assigns in PEQA (see Step 7 of Algorithm 1) on the peer-grading performance score can vary across different instructors. It is, therefore, desirable to have a score-assignment function that is robust to any choice of the weight α while retaining the two properties above.

Without incentive concerns, a statistician would suggest a score-assignment function that would minimize the expected squared distance between the assigned score and true score on exam j , conditional on the true bias and variance parameters. Then, those true parameters could be approximated by the estimated bias and variance. In Equation (26) of Appendix B, we show that under the strong distribution-assumptions of Piech et al. (2013), such a score-assignment function on exam j would look like a weighted-aggregator:

⁵Our score-assignment function is unique in this regard. We show this in Theorem 2.

⁶For an integer k and set N , define a k -subset to be a subset containing exactly k elements. Let $\mathcal{F}_{K/2}^i$ be the set of all $(K/2)$ -subsets of the set $N \setminus \{i\}$. Note that $K/2$ is the number of maximum possible co-graders of i assigned to a given non-probe paper under PEQA. The proof of Theorem 1 can be extended (we provide a sketch in the proof) to show that it holds under the weaker 'if' condition of $\max_{i \in N} \max_{A \subset \mathcal{F}_{K/2}^i} \sum_{k \in A} w_{ik} \leq 1$.

$$r_j = \frac{\lambda_0 \times \text{prior mean of all exams} + \sum_i \lambda_i \times \text{de-biased report on exam } j \text{ from grader } i}{\lambda_0 + \sum_i \lambda_i} \quad (9)$$

with the weights λ_i equal to the inverse of grader-variance. In a world where grading precision needs to be incentivized, and there are constraints on how large grading performance scores might be, this is no longer the optimal choice from this class. The following result states why our score-assignment function, that weighs observations by the inverse of the square root of variance, works better.

THEOREM 2 (Uniqueness) *Assume for every $i \in N$, $\exists j \in G^{-1}(i)$ s.t. $w_{ij} > 0$. A peer-grading mechanism $M = \langle G, \mathbf{r}^*, \mathbf{t} \rangle$, where $r_j^* \equiv r_j^{\text{WA}}, \forall j \in N$ (Equation (4)) with any performance score function t satisfies EPRM for every peer-grading performance score weight $\alpha > 0$ and for all realizations of y_j and $\tilde{y}_j^{(i)}, i \in N, j \in P_i$, if and only if $\lambda_i = c_i/\sigma_i$, where $c_i > 0$ is any arbitrary constant independent of σ_i .*

$\sqrt{\hat{\tau}_i} \propto 1/\sigma_i$ ⁷. Therefore, in the class of weighted average score computing function, ISWDM score-assignment function, used by PEQA uniquely (upto constant multipliers) ensures EPRM for flexible performance score weight $\alpha > 0$.

This second result shows why our score-assignment function is special, irrespective of the choice of grading performance scores. For example, the grading performance score need not be equal to marginal contributions to some social reward. For arbitrary performance score functions, if the course-designer wants to use a small weight on the performance scores in her course, the competitiveness of graders must also be utilized to fuel the incentives for precision. For competitiveness to incentivize precision, noisier grading must result in higher expected scores for others, post regrading requests. Theorem 2 shows that under see Assumption 1, this is only guaranteed when the assignment function uses weights equal to the inverse *square-root* of the grader-variance.

At the risk of oversimplification, here is an intuition about how this result works. For the class of weighted average (WA) score-assignment functions, let us consider how the weights affect the post-regrading score.

$$\max \left\{ r_j^{\text{WA}}(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j \right\} = y_j + \max \left\{ \frac{\lambda_0(\mu - y_j) + \sum_{i \in G(j)} \lambda_i (\tilde{y}_j^{(i)} - \hat{b}_i - y_j)}{\lambda_0 + \sum_{i \in G(j)} \lambda_i}, 0 \right\} \quad (10)$$

The first term in the post-regrading score is the true score y_j which is independent of grader i 's actions. We will focus on how grader i 's choices affect the numerator and the denominator of the second term, that is a non-negative number. In the numerator, the $(\tilde{y}_j^{(i)} - \hat{b}_i - y_j)$ terms are approximately a measure of the noise present in the signals that grader i observed for paper j and her probes. This noise depends directly on the variance σ_i^2 of i 's signal distribution. In general, each term $(\tilde{y}_j^{(i)} - \hat{b}_i - y_j)$ in the numerator can be positive or negative, depending on the directions of noise for the different graders. This creates a potential non-systematic relationship between σ_i and the whole fraction, a relationship that also depends on the sign, magnitude and distribution of the noise present in other's signals. But, $\lambda_i = c_i/\sigma_i$ uniquely de-links the product $\lambda_i(\tilde{y}_j^{(i)} - \hat{b}_i - y_j)$ from grader i 's chosen σ_i , for all values of σ_i . And this is true for all her co-graders too. The numerator is now independent of the variance or precision chosen by the graders. This is the first step in establishing a monotonic relation of the second term with σ_i for any i .

⁷The proportionality constants depend on the realized values of y_j and $\tilde{y}_j^{(i)}, j \in P_i$, after division by σ_i

The denominator is the sum of estimated precisions, all positive numbers. In the denominator, $\lambda_i = c_i/\sigma_i$ guarantees that when σ_i increases, the numerator decreases and the whole fraction increases. Thus, noisier grading ends up increasing the post-regrading score $\max \left\{ r_j^{\text{WA}}(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j \right\}$. In the proof, we formalize this intuition, while accounting for what information is available to grader i when she contemplates how her actions affect post-regrading scores.

6 Efficiency Under Costly Effort

In this section, we will show that a modification of the grading performance score \mathbf{t}^* implements the socially optimal level of effort under costly effort. We assume that the graders have a uniform weight for the other-regarding component in their utility ($w_{ij} = w, \forall i, j \in N$), which is a common knowledge.

Costly effort. We assume that all graders face the same effort-cost function c_j while grading paper j . But, grading different papers at the same reliability might require different effort-costs: we allow $c_j \neq c_k$. We assume that graders can observe c_j while they grade a non-probe paper j , but the mechanism designer cannot.

Grader i chooses a reliability vector $\tau_i = (\tau_{ij}, j \in G^{-1}(i))$, where τ_{ij} is the reliability specific to the paper j that i grades. Reliability is bounded above, i.e., $\tau_{ij} \in [0, \bar{\tau}], \forall i, j \in N$. We summarize our assumptions below.

- (i) The cost $c_j : [0, \bar{\tau}] \rightarrow \mathbb{R}_{\geq 0}$ is convex, increasing, and equal for all graders $k \in G(j)$.
- (ii) The course instructor does not know c_j , only the graders of paper j do.
- (iii) The cost of grading paper j for an agent depends only on the reliability chosen by that agent on that paper, i.e., $c_j(\tau_i) = c_j(\tau_{ij})$.

Social utility of grading: For any non-probe paper, we assume that the social planner (e.g., the instructor) cares about two things: (a) the accuracy of the final score (measured by the reward function $R(r_j^*, y_j)$) and (b) the total cost of grader-effort. We also assume that *if the social planner was aware of the cost functions of grading*, she would have recommended a joint strategy profile (τ_i, τ_{-i}) that maximizes some linear combination of the reward and cost factors as the *social utility*. Hence, we write the social utility of grading paper j as

$$E_{y_j} E_{(\tilde{y}_j^{(k)} | y_j) \sim \mathcal{F}(y_j + b_k, 1/\tau_k), k \in G(j)} \left(\beta R(r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j) - \sum_{i \in G(j)} c_j(\tau_i) \right),$$

where $\beta > 0$ determines the relative weight between the two factors. The final social utility is the sum of this over all non-probe papers.

Aligning social and individual incentives: There are three challenges on the way to aligning social and individual incentives. We discuss these below along with their solutions.

- ▷ A class-designer would care about the accuracy in grade-allocation, but why would graders care about the same? PEQA's grading performance score solves this. It forces subjects to internalize accuracy in their decisions by paying each grader their marginal contribution to grading accuracy.
- ▷ Each competitive grader wants lower grades for others as part of her other-regarding utility. Clearly this is not aligned with the social utility of grading and becomes a source of externality. We solve this by suggesting a modified grading performance score (see below), that additionally compensates graders for any potential losses from their other-regarding utility component.

- ▷ The solution to the last bullet-point presents a new challenge. The other-regarding utility component would be different for each grader i , as their reference groups $N \setminus \{i\}$ would be naturally different. Thus they would have to be compensated different amounts. Would this change the ordinal ranking of students in the class? We show how our method solves this in Lemma 1 below.

Modified grading performance score: Let the post-regrading request grade be $g_i = \max\{r_i, y_i\}$. We propose the modified grading performance score

$$\pi_i := t_i + w \sum_{j \in N \setminus \{i\}} (g_j + \pi_j),$$

where t_i is the original PEQA grading performance score. The additional terms in the performance score compensate for the other-regarding externalities. This expression of the modified performance score helps our understanding of how the score compensates other-regarding externalities. But, the simplicity comes at a price: if i and j are co-graders then π_i has been described as a function of π_j and vice-versa! How is the designer supposed to decide the values of π_i and π_j given the interdependency? We show that π_i has an alternative expression that is independent of π_j s.

$$\pi_i = \frac{t_i + w \sum_{j \in N \setminus \{i\}} g_j + w\pi}{1 + w}, \quad (11)$$

$$\text{where, } \pi = \frac{t + w(n-1)g}{1 - w(n-1)}, \text{ and } t = \sum_{i \in N} t_i, \quad g = \sum_{i \in N} g_i, \quad w \neq \frac{1}{n-1}. \quad (12)$$

The game of peer-grading. Accounting for effort-cost, grader i 's utility is

$$u_i = g_i + \pi_i - w \sum_{j \in N \setminus \{i\}} (g_j + \pi_j) - \sum_{j \in G^{-1}(i)} c_j(\tau_i) = g_i + t_i - \sum_{j \in G^{-1}(i)} c_j(\tau_i). \quad (13)$$

The following result shows that it retains the same order of the scores as in the PEQA performance score.

LEMMA 1 (Order Invariance) *The modified PEQA performance score π retains the same order among the students as the original PEQA performance score t .*

Proof: $g_i + \pi_i > g_k + \pi_k \Leftrightarrow g_i + \frac{t_i + w \sum_{j \in N \setminus \{i\}} g_j + w\pi}{1+w} > g_k + \frac{t_k + w \sum_{j \in N \setminus \{k\}} g_j + w\pi}{1+w} \Leftrightarrow \frac{g_i + t_i + w(g+\pi)}{1+w} > \frac{g_k + t_k + w(g+\pi)}{1+w} \Leftrightarrow g_i + t_i > g_k + t_k. \quad \blacksquare$

The following result shows that the modified grading performance score \mathbf{t}^* implements the social optimal level of effort at every paper j .

THEOREM 3 *For every paper $j \in N$, if the course designer uses the modified grading score π_i and sets $\alpha = \frac{\beta}{p_{NP}}$, every maxima of the expected social utility is a Pure Strategy Nash Equilibrium (PSNE) of the induced game among the graders of the paper.*

Proof: The modified PEQA already compensates for the other-regarding component and makes it inconsequential.

The residual performance score t_i is the sum of performance scores t_i^j from each paper $j \in G^{-1}(i)$. Now, $t_i^j = \alpha(W_j^* - W_j^{(-i)*})$, and $W_j^{(-i)*}$ does not depend on i 's reliability τ_i . Hence the part of the i 's utility expression that depends on τ_i and is related to grading paper j is:

$$\alpha W_j^* - c_j(\tau_i) = \alpha R(r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j) - c_j(\tau_i) \quad (14)$$

$\tilde{y}_j^{G(j)}$ is the profile of all reports collected from $G(j)$. Thus it depends on the bias and reliability of co-graders, which is chosen strategically and simultaneously. Grader i is uncertain whether paper j is a probe versus a non-probe paper, and only the latter provides a performance score. Since the proportion p_{NP} of non-probes is announced by the designer, i assigns a probability p_{NP} to any paper being a non-probe. For any choice of bias and reliability by all the graders on paper j , the expected reward to the designer is denoted by $E_{y_j} E_{(\tilde{y}_j^{(k)} | y_j) \sim \mathcal{F}(y_j + b_k, 1/\tau_k), k \in G(j)} R(r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}), y_j)$ ⁸. From the analysis of Theorem 1, we know that this function is independent of bias under PEQA. Therefore, for simplicity, we assume that every grader strategizes only on her reliability. To emphasize the strategic and simultaneous choice of reliability, we rewrite $R()$ using the shorthand $\bar{R}(\tau_i, \tau_{-i})$.

Hence, for any reliability profile chosen by the set of graders on paper j , the part of the expected utility of grader i that depends on τ_i is

$$U_i^j(\tau_i, \tau_{-i}) = \alpha \cdot p_{\text{NP}} \cdot \bar{R}(\tau_i, \tau_{-i}) - c_j(\tau_i). \quad (15)$$

When $\alpha = \frac{\beta}{p_{\text{NP}}}$, grader i 's relevant utility becomes $U_i^j(\tau_i, \tau_{-i}) = \beta \bar{R}(\tau_i, \tau_{-i}) - c_j(\tau_i)$. Let the social optimal be obtained at $\tau^* = (\tau_k^*, k \in G(j))$. Then, it must be that $\beta \bar{R}(\tau_i^*, \tau_{-i}^*) - c_j(\tau_i^*) \geq \beta \bar{R}(\tau_i, \tau_{-i}^*) - c_j(\tau_i)$ for all reliabilities τ_i of player $i \in G(j)$, because, if any grader i could increase her expected utility by choosing any other τ_i , then it is easy to show that it will contradict the optimality of the social utility at τ^* . Thus, clearly if all except i choose τ_{-i}^* , player i cannot do any better than choosing τ_i^* . Thus, $\tau^* = (\tau_k^*, k \in G(j))$ is a PSNE of this game. ■

7 Experimental study

We ran experiments for both PEQA and median mechanisms. Our experimental study serves two goals. We observe the practical trade-off between the theoretical desirability of PEQA against the simplicity of the median mechanism. We also test two of our modeling assumptions: if bias and reliability were identical in probes and non-probes, and, if competitive ($w_{ij} \geq 0$) preferences are a good approximation of the peer-grader behavior.

7.1 Experimental Design

We ran two experimental sessions: One with the median scoring mechanism (27 students), another with the PEQA mechanism (42 students). We recruited subjects through two open-calls to undergraduate students enrolled in a computing course (Prog101). The open calls did not contain any particulars of the two sessions. Every student who signed up for participation was assigned to one of the two sessions.

In many classes, instructors grade on a curve: final numerical scores are converted to letter-grades (A to D) based on relative rankings. Grading on a curve creates a competitive classroom-environment that we wanted to replicate. We told participants that their total score is the sum of their peer-evaluated score and grading performance score. We paid subjects by the relative ranking of total scores in the class, in both sessions. The students who ranked in the first quartile of the total scores received M 650,⁹ the next three quartiles received M 450, M 250, and

⁸The distributions of these random variables are common knowledge of the graders.

⁹M = Indian Rupee (₹), a difference of M 200 is significant for a student.

M 50 respectively. They also received a show-up fee of M 50, irrespective of their total score. The monetary payments were placeholders for grades A to D in a class that grades on a curve: high relative performance resulted in high rewards.¹⁰

In the median mechanism session, all subjects were assigned an identical grading performance score of zero. The total-score ranking was identical to their peer-evaluated-score ranking. Thus, a subject could decrease others' scores on the peer-evaluation task to increase her relative ranking and payment. The median mechanism is used in practice for peer-grading MOOCs, e.g., Coursera uses this mechanism across multiple courses (Coursera, 2021).

The PEQA session used the PEQA assignment and grading performance scores. Thus, manipulation on the peer-evaluation task risked getting a lower performance and total score, which would result in a lower payment.

The instructions and incentive-scheme, included in Appendix D, were explained in detail before each of the sessions began. In both sessions, we used numerical examples in our explanation. For PEQA we showed the relation between performance score and grading precision through a graph and verbally summarized the monotonic relationship.

We conducted both sessions during the weekly Prog101 labs, that happen in a large computer lab. Our study lies at the intersection of Lab and Field experiments. We are interested in peer-grading behavior and the students are our population of interest. In this study, we observed our population of interest in their *naturally occurring environments*, like in Field experiments.

In both sessions, we asked subjects to peer-grade the same weekly class-quiz. We partitioned each quiz into three sub-quizzes (by treating one(two) question(s) of the quiz as a sub-quiz¹¹), and divided each session into *three* rounds. In every round, the subjects were asked to peer-grade five sub-quizzes (each corresponding to one of five of her anonymous peers). At the end of each round, subjects saw: (a) how peers had evaluated her performance on the sub-quiz, (b) her assigned score (median-scoring or PEQA), and (c) how her co-graders that round had evaluated the sub-quiz.

Within every sub-quiz, some (and not all) of the questions were 'regradable'. The students could raise a regrading request for only those questions at the end of the session. In the PEQA sessions, only the regradable questions were incentivized by the grading performance score. The non-regradable questions used the same assignment function but did not have any grading performance score.

We also graded all the papers ourselves (the instructor graded all of them), and we considered these scores to be the *true scores*. The difference between mechanism assigned scores and true scores is a measure of the quality of these mechanisms.

7.2 Hypotheses and results

Bias is the difference between the true score and the peer-assigned score. It measures the average direction and magnitude of manipulation. PEQA assumes that bias is zero or positive: subjects generally do not manipulate grades upwards (i.e., do not collude). Our first hypothesis builds on this assumption.

Hypothesis 1 *Grade-manipulation is not collusive.*

Our second hypothesis suggests that bias should be higher in the last round for two reasons. First, most repeated interactions have an end-game effect: selfish behavior unravels when no future interactions remain. Second, subjects who have experienced grade-manipulation by others might retaliate as a punishment or reciprocal strategy in the later rounds of the treatment.

Hypothesis 2 *Grade-manipulation or bias peaks in the last round of the game.*

¹⁰The ethics committee did not allow us to use university grades in the Prog101 class as incentives.

¹¹The quiz had more than three questions.

	Subjects	Round	Total grade	Avg bias (% of Total grade)	
				regradable	non-regradable
Median	27	1	1+1	-0.4%	-0.6%
				(-4%,3.2%)	(-3.1%,1.9%)
		2	2+2	1.7%	1.2%
				(-1%,4.4%)	(-0.8%,3.1%)
		3	2+2	16.6%	16.4%
		(12%,21.2%)	(12%,20.7%)		

Table 1: Average bias from 3 rounds grading under the Median mechanism. We report the 95% confidence intervals below the averages.

	Subjects	Round	Total grade	Avg bias (% of Total grade)	
				regradable	non-regradable
PEQA	42	1	1+1	-0.6%	-0.5%
				(-2%,1%)	(-1.9%,0.9%)
		2	2+2	0.6%	-0.3%
				(0%,1.2%)	(-1.2%,0.6%)
		3	2+2	15.8%	15%
		(12.3%,19.2%)	(11.2%,18%)		

Table 2: Average bias from 3 rounds grading under the PEQA mechanism. We report the 95% confidence intervals under the averages.

In Tables 1 and 2, we summarize the bias in individual grading behavior in the three rounds of both treatments. To compare across questions and rounds, we normalize bias by the total score on the relevant question.

Each round every subject graded a regradable and a non-regradable question. The average bias is statistically identical to zero for the first two rounds, and significantly positive in the third round. This is true for both the regradable and non-regradable questions. Thus, the bias is either zero, or positive, and we cannot reject Hypothesis 1. The average bias is also significantly higher in the third round, confirming Hypothesis 2. This holds for both regradable and non-regradable questions.

The low bias in grading in the first two rounds parallels the results on honest reporting from the “die-roll in person and report” studies (Mazar et al., 2008; Fischbacher and Föllmi-Heusi, 2013). In these studies, subjects roll a die privately, self-report the outcome, and get paid based on the report. Fischbacher and Föllmi-Heusi (2013) report that only 20% of people lie to the fullest extent, 39% choose to be honest, and a sizable proportion cheats only marginally. Lying aversion (Dufwenberg and Dufwenberg, 2018), caring about lie-credibility, and a notion of self-concept maintenance (Mazar et al., 2008) are potential reasons for why people do not lie completely even under full anonymity.

How do the two mechanisms perform? The median assignment rule, due to its robustness to outliers, is immune to insincere grading as long as only a minority of graders are insincere. PEQA is bias invariant (EPBI), incentivizes effort, and should outperform the Median mechanism. We use the precision of the mechanism-assigned scores as a metric of relative performance. Given subjects graded most insincerely in the third round, we use this round to test Hypothesis 3.

Hypothesis 3 *In the presence of strategic manipulation, the final score assigned under PEQA should be closer to the true scores, than that assigned under median-scoring.*

In Table 3, we present the means of fractional-difference and squared fractional-difference between the mechanism-assigned score and true score. Thus, for the former we calculate $d_j =$

(true score_{*j*} – mechanism assigned score_{*j*})/total score_{*j*} on student *j*’s third-round sub-quiz, and then take the average over all *j*. Similarly, the latter is the average of d_j^2 . We find the true score on an exam by grading it ourselves.

The average difference between true and mechanism assigned scores is 14.8% under Median and only -1.2% under PEQA. The negative sign indicates that PEQA assigned slightly higher scores than the true score. Both difference and squared difference are significantly smaller under PEQA. Under the median mechanism, the difference and squared difference were equal because d_j almost always took values of 0 or 1.

	Median	PEQA
Mean Difference	14.8%	-1.2%
Mean Squared Difference	14.8%	0.6%
<i>N</i>	27	42

Table 3: Difference and squared-difference

The median mechanism assigned lower than true grade (assigned a 1 instead of a 2) for 15% (4 out of 27) of the sub-quizzes. In comparison, the PEQA mechanism was (almost) always point-precise: only one sub-quiz (out of 41) assigned a grade of 0.5 points higher. The number of regrading requests in the median and PEQA sessions were 4/27 and 3/41 respectively, a difference that is statistically significant.

One of the crucial assumptions of PEQA was that bias and noise are invariant across probes and non-probes.

Hypothesis 4 *Bias and noise are identical in probes and non-probes.*

We tested if bias was different across probe and non-probe questions from the PEQA session. We pooled across three rounds to maximize power. Our statistical tests fail to reject Hypothesis 4. In a t-test, we could not reject the equality of average bias across regradable probe and non-probe questions: the p-value was 0.19. The p-value was 0.16 when we ran the same test for the non-regradable questions.

We also tested if the mean squared deviation (noise) was different across probe and non-probe questions from the PEQA session. We could not reject the equality of noise across probe and non-probe questions. The corresponding p-values were 0.86 and 0.31 respectively for the regradable and non-regradable questions.

8 Conclusions

We introduce a new mechanism, PEQA, that uses an assignment rule and grading performance scores to incentivize competitive graders. The rule and the performance score guarantee unbiased grades. They also guarantee that any grader’s utility increases monotonically with her grading precision, irrespective of her competitiveness and how her co-graders act. Our assignment rule is unique in its class to satisfy this utility-precision monotonicity while allowing flexibility in how large performance scores need to be. When grading is costly, a special version of PEQA implements the socially optimal effort-choices in an equilibrium of the peer-evaluation game among co-graders. Finally, in our classroom experiments, PEQA outperforms the popular median mechanism.

Acknowledgments

This work has been supported by the Indian Institute of Technology Kanpur under the grant number 2017198. We would like to thank the Institutional Ethics Committee (IEC) of the Indian Institute of Technology Kanpur for providing us with the opportunity to run the human subject study with the students of the institute via the IEC Communication Number: IITK/IEC/2020-21/II/30. We thank the student participants of the studies and the survey. We also thank the participants of the Workshop on Economic Design (WED 2019) held at the Indian Statistical Institute Delhi in December 2019.

References

- Noga Alon, Felix Fischer, Ariel Procaccia, and Moshe Tennenholtz. 2011. Sum of us: Strategyproof selection from the selectors. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*. 101–110.
- Yang Cai, Constantinos Daskalakis, and Christos Papadimitriou. 2015. Optimum statistical estimation with strategic data sources. In *Conference on Learning Theory*. 280–296.
- Juan Miguel Campanario. 1998. Peer review for journals as it stands today - Part 1. *Science communication* 19, 3 (1998), 181–211.
- Ioannis Caragiannis, George A Krimpas, and Alexandros A Voudouris. 2015. Aggregating partial rankings with applications to peer grading in massive online open courses. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. International Foundation for Autonomous Agents and Multiagent Systems, 675–683.
- Kwangsu Cho and Christian D Schunn. 2007. Scaffolded writing and rewriting in the discipline: A web-based reciprocal peer review system. *Computers & Education* 48, 3 (2007), 409–426.
- E. H. Clarke. 1971. Multipart Pricing of Public Goods. *Public Choice* 11 (1971), 17–33.
- Coursera. accessed Feb 2021. Online. (accessed Feb 2021). <https://learner.coursera.help/hc/en-us/articles/208279946-See-feedback-and-grades-for-peer-reviewed-assignments>
- Anirban Dasgupta and Arpita Ghosh. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*. ACM, 319–330.
- Luca De Alfaro and Michael Shavlovsky. 2014. CrowdGrader: A tool for crowdsourcing the evaluation of homework assignments. In *Proceedings of the 45th ACM technical symposium on Computer science education*. ACM, 415–420.
- Martin Dufwenberg and Martin A Dufwenberg. 2018. Lies in disguise—A theoretical analysis of cheating. *Journal of Economic Theory* 175 (2018), 248–264.
- Boi Faltings, Jason Jingshi Li, and Radu Jurca. 2012. Eliciting truthful measurements from a community of sensors. In *Internet of Things (IOT), 2012 3rd International Conference on the*. IEEE, 47–54.
- Urs Fischbacher and Franziska Föllmi-Heusi. 2013. Lies in disguise—an experimental study on cheating. *Journal of the European Economic Association* 11, 3 (2013), 525–547.
- Alice Gao, James R Wright, and Kevin Leyton-Brown. 2016. Incentivizing evaluation via limited access to ground truth: Peer-prediction makes things worse. *arXiv preprint arXiv:1606.07042* (2016).

- T. Groves. 1973. Incentives in Teams. *Econometrica* 41 (1973), 617–631.
- John Hamer, Kenneth TK Ma, and Hugh HF Kwong. 2005. A method of automatic grade calibration in peer assessment. In *Proceedings of the 7th Australasian conference on Computing education-Volume 42*. Australian Computer Society, Inc., 67–72.
- Ron Holzman and Hervé Moulin. 2013. Impartial nominations for a prize. *Econometrica* 81, 1 (2013), 173–196.
- Radu Jurca and Boi Faltings. 2005. Enforcing truthful strategies in incentive compatible reputation mechanisms. In *International Workshop on Internet and Network Economics*. Springer, 268–277.
- Radu Jurca and Boi Faltings. 2009. Mechanisms for making crowds truthful. *Journal of Artificial Intelligence Research* 34 (2009), 209–253.
- Vijay Kamble, Nihar Shah, David Marn, Abhay Parekh, and Kannan Ramachandran. 2015. Truth serums for massively crowdsourced evaluation tasks. *arXiv preprint arXiv:1507.07045* (2015).
- Chinmay E Kulkarni, Richard Socher, Michael S Bernstein, and Scott R Klemmer. 2014. Scaling short-answer grading by combining peer assessment with algorithmic scoring. In *Proceedings of the first ACM conference on Learning@ scale conference*. ACM, 99–108.
- Nina Mazar, On Amir, and Dan Ariely. 2008. The dishonesty of honest people: A theory of self-concept maintenance. *Journal of marketing research* 45, 6 (2008), 633–644.
- Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51, 9 (2005), 1359–1373.
- Dwayne E Paré and Steve Joordens. 2008. Peering into large lectures: examining peer and expert mark agreement using peerScholar, an online peer assessment tool. *Journal of Computer Assisted Learning* 24, 6 (2008), 526–540.
- Chris Piech, Jonathan Huang, Zhenghao Chen, Chuong Do, Andrew Ng, and Daphne Koller. 2013. Tuned models of peer assessment in MOOCs. *arXiv preprint arXiv:1307.2579* (2013).
- Dražen Prelec. 2004. A Bayesian truth serum for subjective data. *science* 306, 5695 (2004), 462–466.
- Goran Radanovic and Boi Faltings. 2015. Incentives for subjective evaluations with private beliefs. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI 15)*. 1014–1020.
- Karthik Raman and Thorsten Joachims. 2014. Methods for ordinal peer grading. In *Proceedings of the 20th ACM SIGKDD international conference on Knowledge discovery and data mining*. ACM, 1037–1046.
- Philip M Sadler and Eddie Good. 2006. The impact of self-and peer-grading on student learning. *Educational assessment* 11, 1 (2006), 1–31.
- Nihar B Shah, Joseph K Bradley, Abhay Parekh, Martin Wainwright, and Kannan Ramchandran. 2013. A case for ordinal peer-evaluation in MOOCs. In *NIPS Workshop on Data Driven Education*. 1–8.
- Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C Parkes. 2016. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 179–196.

- Brent Strong, Mark Davis, and Val Hawks. 2004. Self-grading in large general education classes: A case study. *College Teaching* 52, 2 (2004), 52–57.
- W. Vickrey. 1961. Counter Speculation, Auctions, and Competitive Sealed Tenders. *Journal of Finance* 16, 1 (1961), 8–37.
- Bo Waggoner and Yiling Chen. 2014. Output agreement mechanisms and common knowledge. In *Second AAAI Conference on Human Computation and Crowdsourcing*.
- Jens Witkowski, Yoram Bachrach, Peter Key, and David Christopher Parkes. 2013. Dwelling on the negative: Incentivizing effort in peer prediction. In *First AAAI Conference on Human Computation and Crowdsourcing*.
- Jens Witkowski and David C Parkes. 2013. Learning the prior in minimal peer prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*, Vol. 14.
- James R Wright, Chris Thornton, and Kevin Leyton-Brown. 2015. Mechanical ta: Partially automated high-stakes peer grading. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*. ACM, 96–101.

Appendices

A Omitted Proofs

A.1 Proof of Theorem 1

By Assumption 1, the student knows her y_j perfectly and if $r_j^* \geq y_j$, she does not raise a regrading request. PEQA will assume r_j^* to be the true score and design the peer-grading performance score accordingly when there is no regrading request. The student asks for regrading *only if* $r_j^* < y_j$. The utility of grader i after the regrading requests have been addressed is (we have omitted the arguments of the functions in Equation (6) where it is understood) therefore

$$u_i(\cdot) = \max\{r_i^*(\cdot), y_i\} + t_i - \left(\sum_{j \in G^{-1}(i)} w_{ij} \max\{r_j^*(\cdot), y_j\} + \sum_{k \in CG_i \setminus \{i\}} w_{ik} t_k \right) - \phi(\cdot) \quad (16)$$

We decomposed the utility expression to show the terms that are affected by the choices of b_i, τ_i of student i . They are (a) the exam scores of the papers graded by i (first term in the parentheses), and (b) the peer-grading performance score of the co-graders of i (second term in the parentheses). The function ϕ is the remaining part of u_i that is independent of b_i, τ_i .

We prove that PEQA is EPBI and EPRM in *four* steps. First, we observe that the first term on the RHS is independent of the values of b_i and τ_i . In the second step, we show that each summand $\max\{r_j^*(\cdot), y_j\}$ in the first summation term is independent of b_i and decreasing in τ_i . The third step shows that t_i is independent of b_i and increasing in τ_i , and the fourth step shows that this conclusion is true even for $t_i - \sum_{k \in CG_i \setminus \{i\}} w_{ik} t_k$ for the sufficient condition of the theorem.

Step 1: $\max\{r_i^*(\cdot), y_i\}$ is independent of the values of b_i and τ_i . This is obvious since student i does not grade her own paper and hence she has no control on the grade given by PEQA on her paper.

Step 2: The expected value of $\max\{r_j^*(\cdot), y_j\}$ is independent of b_i and increasing in σ_i . Recall that the score-assignment function for PEQA is ISWDM (Definition 1)

$$r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}) = \frac{\sqrt{\gamma}\mu + \sum_{i \in G(j)} \sqrt{\hat{\tau}_i}(\tilde{y}_j^{(i)} - \hat{b}_i)}{\sqrt{\gamma} + \sum_{i \in G(j)} \sqrt{\hat{\tau}_i}}.$$

The final grade after regrading is

$$\max\{r_j^*(\cdot), y_j\} = \max\{r_j^*(\cdot) - y_j, 0\} + y_j. \quad (17)$$

Grader i 's estimated bias is given by $\hat{b}_i = \frac{\sum_{k \in P_i} (\tilde{y}_j^{(i)} - y_j)}{x}$, where $x = |P_i|$. In PEQA, we use the same number $K/2$ as $|P_i|$, for all i . Hence, $x = K/2$, is a constant in our analysis.

Given our model of peer-reports, $\tilde{y}_j^{(i)} = y_j + b_i + n_{ij}$, where $n_{ij} \sim \mathcal{F}(0, 1/\tau_i)$ is a noise term. Hence, it is easy to show that $\hat{b}_i = b_i + \frac{\sum_{k \in P_i} n_{ik}}{x}$ and $\frac{1}{\hat{\tau}_i} = \hat{\sigma}_i^2 = \frac{\sum_{k \in P_i} (n_{ik} - \frac{1}{x} \sum n_{ik})^2}{x}$, where $n_{ik} \sim \mathcal{F}(0, \sigma_i^2)$.

Substituting these values we get the expression for

$$r_j^*(\cdot) - y_j = \frac{\sqrt{\gamma}(\mu - y_j) + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}}.$$

Note that $z_j = \sqrt{\gamma}(\mu - y_j)$ is a $\mathcal{F}(0, 1)$ variable, that is independent of all the other variables in the expression. In the following, we take the expectation of the term $\max\{r_j^*(\cdot) - y_j, 0\}$ w.r.t. z_j and show that it is independent of b_i and increasing in $\sigma_i = 1/\sqrt{\tau_i}$, which implies that irrespective of the values of the other graders' biases and reliabilities, it is best for grader i to reduce her σ_i to increase this component of her utility (since the term comes with a negative sign in the utility expression).

$$\begin{aligned}
I_j &= E_{z_j} \max\{r_j^*(\cdot) - y_j, 0\} \\
&= \int_{-\sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})}^{\infty} \frac{z_j + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} f(z_j) dz_j + 0 \\
&= \frac{1}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \int_{-\sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})}^{\infty} (z_j + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})) f(z_j) dz_j \\
&= \frac{1}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \int_0^{\infty} v_j f(v_j - \sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})) dv_j \\
&= \frac{1}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \int_0^{\infty} v_j f(v_j - \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x}) - \sqrt{\hat{\tau}_i} (n_{ij} - \frac{\sum_{k \in P_i} n_{ik}}{x})) dv_j \\
&= \frac{1}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \int_0^{\infty} v_j f(v_j - \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x}) - \frac{(m_{ij} - \frac{\sum_{k \in P_i} m_{ik}}{x}) \sigma_i}{\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum m_{ik})^2}{x} \sigma_i^2}}) dv_j \\
&= \frac{1}{\sqrt{\gamma} + \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l} + 1/\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum m_{ik})^2}{x} \sigma_i^2}} \times \\
&\quad \int_0^{\infty} v_j f(v_j - \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x}) - \frac{(m_{ij} - \frac{\sum_{k \in P_i} m_{ik}}{x})}{\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum m_{ik})^2}{x}}}) dv_j \tag{18}
\end{aligned}$$

In the third equality, we have substituted $v_j = z_j + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l} (n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})$, and in the fifth equality, we substituted $n_{ik} = m_{ik} \cdot \sigma_i$. Since $n_{ik} \sim \mathcal{F}(0, \sigma_i^2)$, we get $m_{ik} \sim \mathcal{F}(0, 1)$. Note that f is the density of a $\mathcal{F}(0, 1)$ random variable. Hence the whole expression within the integral is independent of σ_i . It is easy to see that the pre-multiplied term is increasing in σ_i . Hence, we conclude that the integral I_j is independent of b_i and increasing in $\sigma_i = 1/\sqrt{\tau_i}$.

Step 3: The expected value of t_i^j is independent of b_i and decreasing in σ_i . We assumed in §2 that the reward function is decreasing in the difference $|r_j^* - y_j|$ and the mechanism assigns reward to be zero when $r_j^* > y_j$. Hence, we calculate the condition on y_j when the reward

is non-zero.

$$\begin{aligned}
r_j^*(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)}) \leq y_j &\iff \frac{\sqrt{\gamma}\mu + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}(\tilde{y}_j^{(l)} - \hat{b}_l)}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \leq y_j \\
&\iff \sqrt{\gamma}\mu + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}(\tilde{y}_j^{(l)} - \hat{b}_l) \leq y_j(\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}) \\
&\iff y_j\sqrt{\gamma} \geq \sqrt{\gamma}\mu + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}(\tilde{y}_j^{(l)} - \hat{b}_l - y_j) = \sqrt{\gamma}\mu + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}(n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x}) \\
&\iff y_j \geq \frac{\sqrt{\gamma}\mu + Z + \sqrt{\hat{\tau}_i}(n_{ij} - \frac{\sum_{k \in P_i} n_{ik}}{x})}{\sqrt{\gamma}}, \quad \text{where } Z = \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l}(n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x}) \\
&= \frac{\sqrt{\gamma}\mu + Z}{\sqrt{\gamma}} + \frac{(m_i - \frac{\sum_{k \in P_i} m_{ik}}{x})\sigma_i}{\sqrt{\gamma}\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum m_{ik})^2}{x}}\sigma_i^2}
\end{aligned}$$

Note that the RHS is independent of σ_i . Hence the limits of the integral where the reward R is non-zero is also independent of σ_i .

By definition, the $W_j^{(-i)*}$ component of the performance score is independent of bias and reliability of grader i . Hence, we only consider the first component which is dependent on the bias and reliability of grader i . We will consider the integral only w.r.t. y_j to compute t_i^j and we just showed that the limits of this integral is independent of σ_i . Hence, if we show that the reward function $R(r_j^*, y_j)$ is independent of b_i and decreasing in σ_i , then we are done. Consider the argument of the reward function

$$\begin{aligned}
r_j^* - y_j &= \frac{\sqrt{\gamma}(\mu - y_j) + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}(n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x})}{\sqrt{\gamma} + \sum_{l \in G(j)} \sqrt{\hat{\tau}_l}} \\
&= \frac{[\sqrt{\gamma}(\mu - y_j) + \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l}(\tilde{y}_j^{(l)} - \hat{b}_l - y_j)]\sqrt{\frac{\sum_{k \in P_i} (n_{ik} - \frac{1}{x} \sum_{l \in P_i} n_{il})^2}{x}} + (n_j - \frac{1}{x} \sum_{l \in P_i} n_{il})}{(\sqrt{\gamma} + \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l})\sqrt{\frac{\sum_{k \in P_i} (n_{ik} - \frac{1}{x} \sum_{l \in P_i} n_{il})^2}{x}} + 1}} \\
&= \frac{Z_{-i}\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum_{l \in P_i} m_{il})^2}{x}} \cdot \sigma_i + \sigma_i \cdot (m_j - \frac{1}{x} \sum_{l \in P_i} m_{il})}{X_{-i}\sqrt{\frac{\sum_{k \in P_i} (m_{ik} - \frac{1}{x} \sum_{l \in P_i} m_{il})^2}{x}} \cdot \sigma_i + 1}} \tag{19}
\end{aligned}$$

In the last equality, we substituted $Z_{-i} = [\sqrt{\gamma}(\mu - y_j) + \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l}(\tilde{y}_j^{(l)} - \hat{b}_l - y_j)]$ and $X_{-i} = (\sqrt{\gamma} + \sum_{l \in G(j) \setminus \{i\}} \sqrt{\hat{\tau}_l})$. As before, we substituted $n_{ik} = m_{ik} \cdot \sigma_i$. Since $n_{ik} \sim \mathcal{F}(0, \sigma_i^2)$, we get $m_{ik} \sim \mathcal{F}(0, 1)$. We see that the absolute value of the above expression is independent of b_i and increasing in σ_i . Hence $R(r_j^*, y_j)$ is independent of b_i and decreasing in σ_i .

Step 4: $t_i^j - \sum_{k \in CG_i^j \setminus \{i\}} w_{ik} t_k^j$ is independent of b_i and decreasing in σ_i for $\sum_{k \in N \setminus \{i\}} w_{ik} \leq$

1. First, we show that $t_i^j - t_k^j$ is independent of b_i and decreasing in σ_i . This is because W_j^* cancels and this difference reduces to $W_j^{(-k)*} - W_j^{(-i)*}$. The second term is independent of b_i and σ_i . The first term is independent of b_i and decreasing in σ_i by the same argument as step 3, with the set of graders reduced to $N \setminus \{k\}$.

Observe that, in the utility of grader i , the difference in these two performance score terms

appear as follows.

$$t_i - w_{ik} \cdot \sum_{k \in CG_i \setminus \{i\}} t_k = \sum_{j \in NP_i} \left(t_i^j - w_{ik} \cdot \sum_{k \in CG_i^j \setminus \{i\}} t_k^j \right).$$

Consider the terms in the parentheses on the RHS.

$$t_i^j - w_{ik} \cdot \sum_{k \in CG_i^j \setminus \{i\}} t_k^j = \sum_{k \in CG_i^j \setminus \{i\}} w_{ik} \cdot (t_i^j - t_k^j) + (1 - \sum_{k \in CG_i^j \setminus \{i\}} w_{ik}) t_i^j.$$

Both terms in the RHS is independent of b_i and decreasing in σ_i as we have already shown and since $\sum_{k \in CG_i^j \setminus \{i\}} w_{ik} \leq \sum_{k \in N \setminus \{i\}} w_{ik} \leq 1$. (Note that the first inequality can also be written as $\leq \max_{i \in N} \max_{A \subset \mathcal{F}^i_{K/2}} \sum_{k \in A} w_{ik}$, since $|CG_i^j| \leq K/2 + 1$, which proves the other version of the theorem with a weaker sufficient condition.)

Combining all steps, we have shown that the expected utility of grader i , where the expectation is taken only w.r.t. her true score is independent of b_i and decreasing in σ_i . Hence these two properties hold for any choice of actions by the other graders. Hence we have proved that PEQA is EPBI and EPRM.

A.2 Proof of Theorem 2

Consider the utility expression for agent i with peer-grading performance score weight α

$$\begin{aligned} u_i(\cdot) &= \max\{r_i^*(\cdot), y_i\} + \alpha t_i - \sum_{j \in G^{-1}(i)} w_{ij} \cdot \max\{r_j^*(\cdot), y_j\} - \alpha \sum_{k \in CG_i \setminus \{i\}} w_{ik} \cdot t_k \\ &= \max\{r_i^*(\cdot), y_i\} - \sum_{j \in G^{-1}(i)} w_{ij} \cdot \max\{r_j^*(\cdot), y_j\} + \alpha \left(t_i - \sum_{k \in CG_i \setminus \{i\}} w_{ik} \cdot t_k \right) \end{aligned} \quad (20)$$

As before, the first term on the RHS of the above expression is independent of σ_i . We will show that for all realizations of the random variables $m_{ik}, n_{\ell k}, v_j, \hat{\tau}_\ell$, where $\ell \in G(j) \setminus \{i\}, j \in G^{-1}(i), k \in P_i$, and α , the utility is monotone decreasing in σ_i if and only if $\lambda_i(\sigma_i) = c_i/\sigma_i$. For brevity of notation, we will use λ_i to denote the function where the argument is clear from the context.

We will show that the second term is monotonically decreasing for all realizations of the random variables if and only if $\lambda_i(\sigma_i) = c_i/\sigma_i$. This will complete the proof, since for any other choice of λ_i , the second term will be non-decreasing (derivative is non-negative) for *some* realization of the random variables. The derivative can be zero only when the weights are independent of σ_i 's. But for that case, the t_i 's will also be independent of σ_i , and hence the third term will also have derivative of zero. When the realizations make the derivative increasing, one can choose α to be small enough such that the decay in the third term is always smaller than the increase in the second term, yielding the overall utility to be increasing in σ_i .

Define the following terms to shorten the forthcoming expressions.

$$K_1 = \lambda_0 + \sum_{l \in G(j) \setminus \{i\}} \lambda_l, \quad K_j = \sum_{l \in G(j) \setminus \{i\}} \lambda_l \left(n_{lj} - \frac{\sum_{k \in P_l} n_{lk}}{x} \right), \quad K_{3j} = m_{ij} - \frac{\sum_{k \in P_i} m_{ik}}{x}.$$

Consider the second term in Equation (20). Using Equation (17), we reduce the expression for paper j in the sum into the difference term and consider its expectation w.r.t. $z_j = \lambda_0(\mu - y_j)$, which is a $\frac{\lambda_0}{\sqrt{\gamma}} \mathcal{F}(0, 1)$ random variable, to get a similar expression like Equation (18) as follows.

We ignore the positive constant $\frac{\lambda_0}{\sqrt{\gamma}}$ as it does not play a role in determining the sign of the variation.

$$I_j = E_{\tilde{z}_j \sim \mathcal{F}(0,1)} \max\{r_j^*(\cdot) - y_j, 0\} = \frac{1}{K_1 + \lambda_i} \times \int_0^\infty v_j f(v_j - K_j - K_{3j}\lambda_i) dv_j \quad (21)$$

To find the change w.r.t. σ_i , we take its partial derivative and find $\frac{\partial I_j}{\partial \sigma_i}$ to be

$$\frac{-K_{3j} \left(\frac{\partial(\lambda_i \cdot \sigma_i)}{\partial \sigma_i} \right) (K_1 + \lambda_i) \cdot \int_0^\infty v_j f'(v_j - K_j - K_{3j}\lambda_i) dv_j - \left(\frac{\partial(\lambda_i)}{\partial \sigma_i} \right) \cdot \int_0^\infty v_j f(v_j - K_j - K_{3j}\lambda_i) dv_j}{(K_1 + \lambda_i)^2} \quad (22)$$

Note that K_1 is positive, while K_j and K_{3j} can take any sign. To ensure that the expression above is positive for *all* values of the realized random variables, it is necessary and sufficient that

$$\frac{\partial(\lambda_i \cdot \sigma_i)}{\partial \sigma_i} = 0, \quad \text{and} \quad \frac{\partial(\lambda_i)}{\partial \sigma_i} < 0. \quad (23)$$

This is because the second integral in the numerator is always positive. Therefore, the above condition is equivalent to $\lambda_i = c_i/\sigma_i$, where $c_i > 0$ is a constant. This concludes the proof.

B Calculation of $r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \hat{\boldsymbol{\theta}}_{G(j)})$ under PG_1 (Piech et al., 2013) model:

Below, we find a score-assignment function $r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}, \hat{\boldsymbol{\theta}}_{G(j)})$ that would maximize a quadratic reward function, i.e., that would minimize the expected squared distance between the assigned score and true score on exam j . We will calculate the expression w.r.t. the true error parameters $\boldsymbol{\theta}$. Then, given $\boldsymbol{\theta}$ is not observed, we will approximate $\boldsymbol{\theta}$ with the estimated value of the same parameters, i.e., $\hat{\boldsymbol{\theta}}$ to find a new expression.

To calculate $r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \boldsymbol{\theta}_{G(j)})$, we first need to calculate the conditional distribution of the true score y_j , $\psi(y_j | \tilde{\mathbf{y}}_j^{G(j)}; \mathbf{b}_{G(j)}, \tau_{G(j)}, \mu, \gamma)$ under the PG_1 (Piech et al., 2013) model, where $\psi(\cdot)$ is the density of the normal distribution with the mean and variance given by that model. For the convenience of the reader, we restate the PG_1 model below.

Model PG_1 (grader bias and reliability) This model puts prior distributions over the latent variables and assumes for example that while an individual grader's bias may be nonzero, the average bias of many graders is zero. Specifically,

$$\begin{aligned} (\text{Reliability}) \quad & \tau_v \sim \mathcal{G}(\alpha_0, \beta_0) \quad \text{for every grader } v, \\ (\text{Bias}) \quad & b_v \sim \mathcal{N}(0, 1/\eta_0) \quad \text{for every grader } v, \\ (\text{True score}) \quad & s_u \sim \mathcal{N}(\mu_0, 1/\gamma_0) \quad \text{for every user } u, \\ (\text{Observed score}) \quad & z_u \sim \mathcal{N}(s_u + b_v, 1/\tau_v) \quad \text{for every observed peer grade } s_u, \end{aligned}$$

where \mathcal{G} and \mathcal{N} refers to the gamma and normal distributions respectively with appropriate hyperparameters. The hyperparameters $\alpha_0, \beta_0, \eta_0, \mu_0, \gamma_0$ are the hyperparameters for the priors over reliabilities, biases, and true scores, respectively.

Hence, the conditional distribution of the true score y_j , $\psi(y_j | \tilde{\mathbf{y}}_j^{G(j)}; \mathbf{b}_{G(j)}, \tau_{G(j)}, \mu, \gamma)$ is calculated as follows.

$$\begin{aligned}
\psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}, \mu, \gamma) &= \frac{\psi(y_j; \mu, \gamma) \psi(\tilde{y}_j^{G(j)} | y_j; b_{G(j)}, \tau_{G(j)})}{\int_{y_j} \psi(y_j; \mu, \gamma) \psi(\tilde{y}_j^{G(j)} | y_j; b_{G(j)}, \tau_{G(j)}) dy_j} \\
&\propto \psi(y_j; \mu, \gamma) \psi(\tilde{y}_j^{G(j)} | y_j; b_{G(j)}, \tau_{G(j)}) \\
&\propto \psi(y_j; \mu, \gamma) \prod_{i \in G(j)} \psi(\tilde{y}_j^{(i)} | y_j; b_i, \tau_i) \\
&\propto \exp \left(-\frac{1}{2} \gamma (y_j - \mu)^2 + \sum_{i \in G(j)} \left(-\frac{1}{2} \tau_i (\tilde{y}_j^{(i)} - (y_j + b_i))^2 \right) \right) \\
&\propto \exp \left(-\frac{1}{2} \left[\gamma (y_j - \mu)^2 + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - (y_j + b_i))^2 \right] \right)
\end{aligned}$$

The expression inside the exponent is quadratic. We consider the exponent as follows.

$$\begin{aligned}
&\gamma (y_j - \mu)^2 + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - (y_j + b_i))^2 \\
&= \text{const.} + \gamma (y_j^2 - 2y_j \mu) + \sum_{i \in G(j)} \tau_i \left((y_j + b_i)^2 - 2\tilde{y}_j^{(i)} (y_j + b_i) \right) \\
&= \text{const.} + \left(\gamma + \sum_{i \in G(j)} \tau_i \right) y_j^2 - 2 \left(\gamma \mu + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - b_i) \right) y_j, \\
&= \text{const.} + R \left(y_j - \frac{1}{R} \left(\gamma \mu + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - b_i) \right) \right)^2 \\
&\quad (\text{where } R = \gamma + \sum_{i \in G(j)} \tau_i)
\end{aligned}$$

Therefore the resultant distribution is Gaussian:

$$\psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}, \mu, \gamma) \sim \mathcal{N} \left(\frac{\gamma \mu + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - b_i)}{\gamma + \sum_{i \in G(j)} \tau_i}, \frac{1}{\gamma + \sum_{i \in G(j)} \tau_i} \right)$$

$$\mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}} [y_j] = \frac{\gamma \mu + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - b_i)}{\gamma + \sum_{i \in G(j)} \tau_i} \quad (24)$$

Now we are in a position to calculate $r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \boldsymbol{\theta}_{G(j)})$. The reward function is $R(x_j, y_j) = -(x_j - y_j)^2$ where x_j is the estimated score and y_j is the true score for paper j . The score-assignment rule *expected risk minimizer* (ERM) is given below.

$$\begin{aligned}
r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \boldsymbol{\theta}_{G(j)}) &= \operatorname{argmax}_{x_j \in S} \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) R(x_j, y_j) dy_j \\
&\text{where } b_i, \tau_i \text{ are the estimated bias and reliabilities } \forall i \in G(j) \\
&= \operatorname{argmax}_{x_j \in S} \left[- \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j)^2 dy_j \right] \\
&= \operatorname{argmin}_{x_j \in S} \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j)^2 dy_j
\end{aligned}$$

Let $g_j(x_j) = \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j)^2 dy_j$. Hence we need to find x_j that minimizes $g_j(x_j)$. The first and second order conditions are given as follows.

$$\begin{aligned}
\frac{\partial g_j(x_j)}{\partial x_j} &= \frac{\partial}{\partial x_j} \left[\int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j)^2 dy_j \right] \\
&= \int_{y_j} \frac{\partial}{\partial x_j} \left[\psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j)^2 dy_j \right] \\
&= 2 \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) (x_j - y_j) dy_j \\
&= 2x_j \int_{y_j} \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) dy_j - 2 \int_{y_j} y_j \psi(y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}) dy_j \\
&= 2x_j - 2\mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}} y_j \\
\frac{\partial g_j(x_j)}{\partial x_j} = 0 &\Leftrightarrow x_j = \mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}} y_j \\
\frac{\partial^2 g_j(x_j)}{\partial x_j^2} &= 2 > 0
\end{aligned}$$

The first and second order conditions show that $x_j = \mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}} y_j$ is a global minima. Hence

$$r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \boldsymbol{\theta}_{G(j)}) = \mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; b_{G(j)}, \tau_{G(j)}} y_j = \frac{\gamma\mu + \sum_{i \in G(j)} \tau_i (\tilde{y}_j^{(i)} - b_i)}{\gamma + \sum_{i \in G(j)} \tau_i}. \quad (25)$$

The last equality follows from Equation (24).

Replacing $\boldsymbol{\theta}$ with the estimated parameters, i.e., $\hat{\boldsymbol{\theta}}$, we get,

$$r_j^{\text{ERM}}(\tilde{\mathbf{y}}_j^{G(j)}; \hat{\boldsymbol{\theta}}_{G(j)}) = \mathbb{E}_{y_j | \tilde{y}_j^{G(j)}; \hat{b}_{G(j)}, \hat{\tau}_{G(j)}} y_j = \frac{\gamma\mu + \sum_{i \in G(j)} \hat{\tau}_i (\tilde{y}_j^{(i)} - \hat{b}_i)}{\gamma + \sum_{i \in G(j)} \hat{\tau}_i}. \quad (26)$$

C Comparison of W_j^{ERM} and W_j^{ISWDM}

Figure 1 shows the sub-optimality, i.e., $W_j^{\text{ERM}} - W_j^{\text{ISWDM}}$, with respect to an increasing symmetric bias and reliability, where symmetric means that every peer-grader has the same bias and reliability. It shows that the sub-optimality is small but insensitive to bias and monotone decreasing in reliability.




Figure 1: Sub-optimality with increasing bias and reliability

D Instructions provided to the human subjects (§7)

The instructions for both the mechanisms were as follows.

D.1 Median Mechanism Instructions

First, please register yourself on: [registration link] and solve the problem¹² therein. The example there would help you understand how your decisions map into your final payments, through the median-mechanism and payment system used in this study. This is a study on peer-grading. You should read the following instructions carefully, as they would help you perform successfully in the study. In this study, each of you will be asked to grade the assignments of five anonymous students in this room. Similarly, your own assignment would be graded by five anonymous students from this room. Your peer-graded marks and the relative ranks in this peer-grading exercise only determine your payment from this session. It will not be used to determine your actual score for your final grade in the course. The assignment score used towards your university grades will be provided to you by the instructor (i.e., tutors or myself) later.

Would I know whose exam papers I might be grading / correcting? You would not have this information. We will take maximal precautions to make sure that the grader or the assignment-owner’s identities are anonymous to each other during and after this session. Further you would also not know which other four participants are grading the same papers as you. Thus, this procedure is double-blind. We will provide you a solution manual to help you in the grading process. Follow the explanation of the questions and correct answers presented before the study. Please be respectful and encouraging in the grading process. Scores should reflect the learner’s understanding of the assignment and points should not be deducted for difficulties with language or differences in opinion or for using a different but correct methodology.

How are the final grades on my own assignment decided? All five peer-graders independently assign you grades on all of the questions (there are 5 in total, all worth 2 points). Then for each question-part your final grade is the median of those five grades. For example, if on the second round of peer-grading, the five graders assign you 0, 1, 1.5, 2, and 2 respectively, then

¹²The problem tests the participant’s understanding of median and similar simple techniques.

your final assigned grade on that question would be 1.5. We would calculate your grades on all the questions separately by the above median-method, and then aggregate those median grades from all the questions. For example, if there are five questions and the median grades on the questions are 0, 1, 1.5, 2 and 2 respectively, then the total grade on the assignment is 6.5.

How does one calculate the median of five numbers? Sort the numbers in increasing order and the third highest number would be the median.

Can I dispute my peer-assigned grades? Yes, for certain questions you can, and for others you cannot. In case you think your true grade is different than the grade that has been assigned to you on these questions, you can privately indicate that on a form, that would be sent at the end of the peer-grading and that will immediately notify us. We would then reassign you the grade the Teaching staff had assigned to your assignment previously. This whole process would be completed in a click of a button and you would be shown your updated grade in a matter of seconds. Please note that once a dispute is lodged, your grade would become the Teaching Staff assigned grade irrespective of whether that results in an increase or decrease over your original grade.

How are my payments decided? Every participant would get a show up fee of M 50 for participating in and completing this session. You would also get an additional amount depending on your ranking in the pool of 'n' participants today. The ranking would be done in decreasing order of the final grades assigned to you all on the whole assignment. A ranking of x means that there are (x-1) other people who have a strictly higher grade than you. The additional amount would be equal to M 650 for the top 25% (first quartile) ranked students, M 450 for the next 25% (second quartile) ranked students, M 250 for the third quartile ranked students, and M 50 for the bottom quartile students. If the number of students that scored the same overlaps to two or more different quartiles, then all of them get the average payment of those quartiles. E.g., suppose 7 students got the same marks, and 3 students are in first quartile while 4 are in second quartile – then all 7 get M 600 (average of M 700 and M 500). Hence, in this study, the higher you are in the ranking based on your peers' judgment (and a potential review), higher is your total payment.

How do the grades you submit affect your own payment? The grades you submit obviously do not affect your own grade, because you are never grading your own paper, but they can still affect your own payment. Your grading would potentially affect the grades of others, and that can change the relative rank between you and the person(s) you are grading. For example, when you assign someone a higher grade, that might change the median grade they are assigned, and thus move them to a higher rank than you. Similarly, when you give them a lower grade, it might move them to a relatively lower rank than you. Obviously, both of these scenarios would affect the final payments of both you and the other person, as everyone is paid according to the final rankings.

Time-line for the study in chronological order:

- ▷ Stage 0: The whole assignment to be graded is broken up into 3 small parts, that would be peer-graded in three stages. The total grade from the whole assignment determines your final ranking and payment. At this stage, you are expected to complete the questionnaire successfully.
- ▷ Stage 1: Every one of you peer-grades the first part of the assignment of 5 of your peers. Therefore, for any question you are grading in this stage, you know that 4 other anonymous participants are also grading that question. Also, the first part of your own assignment is also being peer-graded by 5 other participants. One part of these questions will have

options for regrading, while the other part will not (it will be mentioned in the response sheet, but all regrading requests will be collected at the end of stage 3).

- ▷ Feedback Stage 1: For each paper you graded in Stage 1, we will show you the grades assigned by you and the 4 other anonymous graders. We will also show you how part 1 of your own assignment got graded by the assigned graders.
- ▷ Stage 2: Similar to Stage 1, now part 2 of the assignment gets peer-graded. But the papers are now sent to a new random set of peer-graders. One part of these questions will have options for regrading, while the other part will not (it will be mentioned in the response sheet, but all regrading requests will be collected at the end of stage 3).
- ▷ Feedback Stage 2: Feedback of Stage 2 (similar to Stage 1) observed.
- ▷ Stage 3: Similar to Stage 2 (one part has regrading requests, the other does not), now part 3 of the assignment gets peer-graded.
- ▷ Feedback Stage 3: Feedback of Stage 3 (similar to Stages 1 and 2) is sent to all students, along with their tentative total score. You may raise regrading requests for the part that is regradable (as mentioned above). Any regrading requests that are lodged will be acted on. Performance on the whole assignment is aggregated, and the final ranking and payments are sent via email. To finish the study, complete the survey that comes in the last email. Study ends.

Is my data confidential? Yes, your data is completely confidential. Before observing and analyzing the collected data, we would be removing every personal identifier from the data, so that none of the decisions can be traced back to the individual who made the decision.

The first practice example tests you on your understanding of the mechanism how the peer-grading leads to your final grade, rank, and payment. You must complete this practice example with a score of 80% or more (i.e., correctly answer at least 4 questions out of 5). You will get one chance only, so please do this carefully. Failing this, you would be asked to leave this session with a M 20 reward.

Important: Please do not communicate with any other participants during this session. For the grading, open one file at a time, finish grading, submit the grade in the google form and then move on. Please keep seated even if you are done with grading before time. If you have any questions, please raise your hand and one of us will come by to answer your query. Please use your university domain email id throughout this session. Please come remembering your google id/password, since that may be needed for some form filling.

D.2 PEQA Instructions

Before you begin, please register yourself on: [registration link]. Submit the form only once.

This is a study on peer-grading. In this study, each of you will be asked to grade **five** anonymous assignments. Similarly, your own assignment would be graded by a certain number of anonymous students from this room. Your peer-graded marks and your performance in the peer-grading exercise will only determine your payment from this session. It will not be used to determine your actual score for your final grade in the course. The assignment score used towards your university grades will be provided to you by the instructor (i.e., tutors or myself) later.

Would I know whose exam papers I might be grading / correcting? You would not have this information. We will take maximal precautions to make sure that the grader or the assignment-owner’s identities are anonymous to each other during and after this session. Further you would also not know which other four participants are grading the same papers as you. Thus, this procedure is double-blind. We will provide you a solution manual to help you in the grading process. Follow the explanation of the questions and correct answers presented before the study. Please be respectful and encouraging in the grading process. Scores should reflect the learner’s understanding of the assignment and points should not be deducted for difficulties with language or differences in opinion or for using a different but correct methodology.

How are the final grades on my own assignment decided? Your peer-graders independently assign you grades on all of the questions. Then for each question your final grade is decided by running it through a **new mechanism called PEQA (Peer Evaluation with Quality Assurance)**. This is a mechanism which is designed to remove the individual biases in grading, and selectively weight and reward graders by how precise they are (details to follow). We would calculate your grades on all the questions separately by the above method, and then aggregate those grades from all the questions. In each round, you will have some **regradable** and **non-regradable** questions. For the regradable part, you will earn the peer-given score computed through PEQA **and** an additional PEQA reward for grading. For the non-regradable part, you will only receive the peer-given score computed through PEQA, **but no additional reward for grading**.

Can I dispute my peer-assigned grades? Yes, for certain questions you can, and for others you cannot. In case you think your true grade is different than the grade that has been assigned to you on these questions, you can privately indicate that on a form, that would be sent at the end of the peer-grading and that will immediately notify us. We would then reassign you the grade the Teaching staff had assigned to your assignment previously. This whole process would be completed in a click of a button and you would be shown your updated grade in a matter of seconds. Please note that once a dispute is lodged, your grade would become the Teaching Staff assigned grade irrespective of whether that results in an increase or decrease over your original grade.

What is the PEQA mechanism? Let us describe PEQA in short in the following two steps:
Step 1 Probes: Out of the five questions you (a grader) grade, two are randomly assigned to be *probes* (rest three are *non-probes*). On the probe papers, we would directly assign the teaching staff assigned grades and also use the teaching staff assigned grades to get an estimate of your individual *average deviation (or bias)* and *variance* in the assignments you graded. We will do this for all the graders. For a grader who on average, assigns a grade higher than the true-grade, the estimated deviation would be negative, and otherwise would be positive.
Step 2 Non-Probes: The non-probes would be graded using the information from, (i) the assigned grades of all the graders, and (ii) the estimated average deviation (or bias) and variance of grading by peer-graders in Step 1. The assigned scores would be “de-biased” using the information in 2.

Here is a numerical example that goes through these two steps. Suppose on the five questions you graded, the first two questions are randomly assigned as probes (this is for illustration only, the actual probes will be interspersed and not the first two, and you won’t know which are the probes).

Paper	Status	Score you assigned (A)	True Score (B)	Deviation (A-B)	Bias=Avg of Deviation
1	<i>Probe</i>	3	3	3-3=0	$\frac{0+(-.5)}{2} = -.25$
2	<i>Probe</i>	2.5	2	2-2.5=-.5	
3		3.5			
4		4			
5		2			

On the probe questions, your evaluation would be compared with the evaluation done by the course instructors (True score), to calculate an average deviation in your grading. We would then use this to calculate the variance of your deviation.

Paper	Status	Score you assigned	True Score	Deviation	Bias=Avg Deviation	Variance of Deviation
1	<i>Probe</i>	3	3	3-3=0	-.25	$\frac{(0+.25)^2+(-.5+.25)^2}{2} = .0625$
2	<i>Probe</i>	2.5	2	2-2.5=-.5		
3		3.5				
4		4				
5		2				

Suppose the (bias,variance) pairs of the other two graders, who are also grading question 4, are (.25, .05) and (-.5,.2) respectively. Suppose the scores they had assigned to the same Q4 was 3 and 2 respectively, while you have given 4 to that question.

Then, the final grade on Q4 (a typical non-probe question) would be calculated as (k_1 and k_2 are some appropriately chosen constants)

$$\text{assigned_score} = \frac{k_1 + \frac{1}{\sqrt{.0625}}(4 + (-.25)) + \frac{1}{\sqrt{.05}}(3 + .25) + \frac{1}{\sqrt{.2}}(2 + .5)}{k_2 + \frac{1}{\sqrt{.0625}} + \frac{1}{\sqrt{.05}} + \frac{1}{\sqrt{.2}}}$$

When we assign the final grade on any non-probe question, we will “**de-bias**” the reports from all the graders by subtracting out the bias, and also selectively over-weight the information from the low-variance graders. We consider the inverse of the square-root of your variance as your **precision of grading**, and use this precision to weight your assigned score on this paper. The accuracy of the mechanism_assigned_score is given by $-(\text{assigned_score}-\text{true_score})^2$.

If you were not one of the graders, and the mechanism only assigned scores using the reports of the other graders,

$$\text{assigned_score_without_you} = \frac{k_1 + \frac{1}{\sqrt{.05}}(3 + .25) + \frac{1}{\sqrt{.2}}(2 + .5)}{k_2 + \frac{1}{\sqrt{.05}} + \frac{1}{\sqrt{.2}}}$$

The new accuracy is $-(\text{assigned_score_without_you}-\text{true_score})^2$. Now, your PEQA performance score from peer-grading question 4 would be calculated as the difference between the accuracy with you, and the accuracy without you. This is intuitively equivalent to you getting paid for your relative contribution in your group towards making the final assigned grade accurate. **The more accurate the assigned score is, when you are included in the group of graders, the higher would be your performance score!**

The PEQA performance score on each question you have graded that is worth x points, is assigned on the scale of $[0, \frac{x}{2}]$. So, in round 1, where each regradable question is worth one point, and you grade a total of 3 non-probe questions, the maximum PEQA performance score you could get is $3 \times 0.5 = 1.5$ and the minimum is 0.

This PEQA grade and performance scores have the following properties:

Bias Invariance: Suppose you had reported grades of $3+x$, $2.5+x$, $3.5+x$, $4+x$, and $2+x$, on all the questions instead, and thus had an individual deviations x points higher than before. This would have no effect on the PEQA performance scores, as it would be de-biased as described above. This is a mathematical property of the mechanism described.

With the new reported grades, your average deviation is changed to $-.25 - x$ from $-.25$. The $+x$ and $-x$ cancel out in the expression of the assigned score, leaving it unchanged.

$$\text{assigned_score} = \frac{k_1 + \frac{1}{\sqrt{.0625}}(4 + x + (-.25 - x)) + \frac{1}{\sqrt{.05}}(3 + .25) + \frac{1}{\sqrt{.2}}(2 + .5)}{k_2 + \frac{1}{\sqrt{.0625}} + \frac{1}{\sqrt{.05}} + \frac{1}{\sqrt{.2}}}$$

Clearly the assigned_score_without_you also cannot change if your bias changes, so your expected PEQA performance score cannot change here!

Precision Monotonicity: For every set of (bias, variance) your co-graders might have, your expected PEQA performance score from the peer-grading task is monotonically increasing in your grading-precision (*precision is the inverse square-root of your variance*). This is a mathematical property that can be easily showed by using calculus and statistics. Thus the more precisely you evaluate a paper in the peer-grading task, (or alternatively the lower your grading variance) the higher your peer-grading score.

Here is a graph that shows how the PEQA performance score changes with the Precision for a grader, who is grading alongside with two graders, one of highest precision and one of lowest precision.

How do you calculate $-(\text{assigned_score} - \text{true_score})^2$? If there is no regrading request, then we would assume that $\text{true_score} = \text{assigned_score}$, and this value is zero. If there are regrading requests, then we would evaluate the paper ourselves and assign the course-instructor assigned score as true_score to calculate the value.

What is my consolidated score? Your consolidated score is the sum of (i) the score on your own assignment (consolidated score from the regradable and non-regradable parts), and (ii) your PEQA performance score (peer-grading score). For example, if the peer-assigned score (computed via PEQA) on your own assignment is x , and your peer-grading score is y , your consolidated score is $x+y$.

How are my payments decided? Every participant would get a show-up fee of M 50 for participating in and completing this session. You would also get an additional amount depending on your ranking in the pool of 'n' participants today, based on the consolidated score. **The ranking would be done in decreasing order of the final grades (i.e., the consolidated score) assigned to you all on the whole assignment.** A ranking of x means that there are $(x-1)$ other people who have a strictly higher consolidated score than you. The additional amount would be equal to M 650 for the top 25% (first quartile) ranked students, M 450 for the next 25% (second quartile) ranked students, M 250 for the third quartile ranked students, and M 50 for the bottom quartile students. If the number of students that scored the same overlaps to two or more different quartiles, then all of them get the average payment of those quartiles. For example, suppose 13 students out of a population of 40 got 10/10, then all 13 get $M (700 \times 10 + 500 \times 3) / 13 = 654$ – the next rank starts from 14. Hence, in this study, the higher is your consolidated score, higher is your total payment.

How do the grades you submit affect your own payment? The grades you submit obviously do not affect your own grade, because you are never grading your own paper, but they can still affect your own payment, in two ways.

1) **By affecting the grade of others:** Your grading could potentially affect the grades of others, *only if the question is chosen as non-probe question*, and consequently that can change the relative rank between you and the person(s) you are grading. For example, when you assign someone a higher/ lower grade on a question that is chosen as a non-probe question, that might change the PEQAassigned quiz score (and thus the consolidated score) they are assigned, and thus affect the relative rankings. But, note that Bias Invariance result described above already tells you that a **different bias would not change the expected quiz scores** of any peers.

2) **By affecting your peer-grading score:** Assigning a higher/ lower score on any question, could change your payments in two ways. If this happened on a question that was chosen as probe, we would be calculating your precision and bias to a different number, and a **lower (respectively higher) precision would result in a lower (respectively higher) marginal impact of your peer-grading reports, and hence, a lower (respectively higher) peer-grading score (and hence lower consolidated score)** for you. If this was a non-probe question instead, then you might be able to change the peer-graded score on that paper, depending on how much weight we assign to your evaluation.

Is my data confidential? Yes, your data is completely confidential. Before observing and analyzing the collected data, we would be removing every personal identifier from the data, so that none of the decisions can be traced back to the individual who made the decision.

You would be given a questionnaire of three questions that tests you on your knowledge of calculation of median. Failure in answering at least two correctly out of those three questions would disqualify you from participation in this study. In this case you would be asked to leave this session with a M 20 reward. Important: Please do not communicate with any other participants during this session. For the grading, open one file at a time, finish grading, submit the grade in the google form and then move on. Please keep seated even if you are done with grading before time. If you have any questions, please raise your hand and one of us will come by to answer your query. Please use your university domain email id throughout this session. Please come remembering your google id/password, since that may be needed for some form filling.