

High-Effort Crowds: Limited Liability via Tournaments

Yichi Zhang and Grant Schoenebeck

School of Information, University of Michigan

April 13, 2023

Abstract

We consider the crowdsourcing setting where, in response to the assigned tasks, agents strategically decide both how much effort to exert (from a continuum) and whether to manipulate their reports. The goal is to design payment mechanisms that (1) satisfy limited liability (all payments are non-negative), (2) reduce the principal’s cost of budget, (3) incentivize effort and (4) incentivize truthful responses. In our framework, the payment mechanism composes a *performance measurement*, which noisily evaluates agents’ effort based on their reports, and a *payment function*, which converts the scores output by the performance measurement to payments.

Previous literature on information elicitation applies a spot-checking or peer prediction mechanism as a performance measurement and then applies a linear payment function to rescale the payments. This method can already achieve some of the above goals in the binary effort setting: either (1), (3) and (4) or (2), (3) and (4). Casting it as a principal-agent problem, we suggest applying a rank-order payment function (tournament) on agents’ scores. In an idealized setting with Gaussian noise, we analytically optimize the rank-order payment function, and identify a sufficient statistic, sensitivity, which serves as a metric for optimizing the performance measurements in terms of eliciting a desired effort with the minimal cost of budget. This helps us obtain objectives (1), (2) and (3) simultaneously. Additionally, we show that adding noise to agents’ scores can preserve the truthfulness of the performance measurements under the non-linear rank-order payment function, which gives us all four objectives.

Our real-data estimated agent-based model experiments reinforce our theoretical results and show that the proposed mechanism can greatly reduce the payment of effort elicitation while preserving the truthfulness of the performance measurement. In addition, we empirically evaluate several commonly considered performance measurements in terms of their sensitivities and strategic robustness.

1 Introduction

Crowdsourcing, on platforms like Amazon Mechanical Turk, suffers from incentive problems. The requesters would like to pay the workers to incentivize high effort/effortful reports. However, workers can increase their payments by spending less time on each task and completing more tasks, which could wastefully spend the requesters’ budgets. At the extreme, which has been extensively studied [4, 28], workers may answer with little effort or even randomly.

Furthermore, in many crowdsourcing settings, it matters not just whether workers exert effort, but how much effort they exert. Lackadaisical workers may provide mediocre-effort work—enough to pass basic checks but still not of a high-quality standard. For example, while labeling tweets for content moderation, people can report whatever is in their minds after reading the first sentence instead of carefully reading the whole tweet, or they can work on a fraction of tweets while skipping the rest. In these and many other cases, effort is not simply binary, but measured on a continuum. Evidence suggests that lackadaisical behaviors may be ubiquitous on crowdsourcing systems. In one study, 46% of Mechanical Turk workers failed at least one of the validity checks which was twice the percentage in student groups [2].

We study the design of payment mechanisms which determine how much the agents should be paid based on their reports on the assigned tasks. Specifically, we focus on practical payment mechanisms that

possess two key properties: *limited liability* (1), which requires the payments to be non-negative; and *budget efficiency* (2), which ensures that the expected payments are not excessively larger than necessary. The former, as always preferred and often required in realistic settings, plays a crucial role in encouraging participation, especially from risk-averse agents; while the latter is important in avoiding extravagant cost of the requester’s budget.

The payment mechanism establishes a game between the agents, where each agent can strategically choose an effort level to maximize her expected utility, such as the difference between the expected payment and the cost of effort. Therefore, another desired property of the payment mechanism is *effort elicitation* (3), which means that the mechanism induces an equilibrium where agents exert a desired level of effort.

Making matters worse, the problem of effort is only one piece of the larger puzzle of strategic behavior. In addition to varying the amount of effort, agents can also manipulate their responses in an attempt to game the mechanism for higher rewards. For example, instead of reporting their true beliefs about the rating of a restaurant, agents may sometimes hedge their scores to align with what they believe to be the most popular answer. However, in any cases, we want agents to truthfully report their information, which is essential for collecting accurate and high-quality data via crowdsourcing. This property of a payment mechanism is called *truthfulness* (4).

In this paper, we ask the following question:

Can we design payment mechanisms that simultaneously satisfy all four objectives?

We provide a positive answer by proposing a two-stage approach for the design of payment mechanisms. First, given all agents’ reports, a *performance measurement* assigns each agent a *performance score* (Fig. 1). For example, both spot-checking mechanisms [12, 31] and peer prediction mechanisms [4, 28, 21] can be used as performance measurements. The former score agents based on their performances on a subset of the tasks with known ground truth, while the latter score each agent according to the correlations in her reports and her peers’ reports which works in the absence of ground truth. Although these performance measurements are primarily proposed to guarantee truthfulness, it is usually argued that the performance score can serve as a noisy measurement of the agent’s effort and thus can be used for effort elicitation. However, in the spot-checking and peer prediction literature, a careful characterization of effort elicitation has only been put forth in the binary effort setting [22, 19].

In the second stage, to achieve limited liability (1) and budget efficiency (2), the requester must carefully choose a *payment function* to convert the performance scores into final payments. For example, a linear payment function pays each agent an affine transformation of her performance score. Is linear payment function the correct term? The key advantage of linear payment functions is that they trivially preserve the truthfulness of the performance measurement, as maximizing expected payment under linear transformations is the same as maximizing expected score. However, as we will see, they are not effective in eliciting effort. Principal-agent literature [6, 7, 14] has been a major contributor to the understanding of effort elicitation and budget efficiency. In principal-agent models, agents are assumed to be strategic only in their choice of effort but not in manipulating their reports. The principal are assumed to observe a noisy measurement of the agent’s effort. The goal is to design a payment function which maps from the noisy observations to payments, so as to maximize the principal’s utility. However, the optimal payment functions are usually non-linear, and thus do not generally preserve the truthfulness of the performance measurements.




Figure 1: Components of a payment mechanism: a performance measurement and a RO-payment function.

Although it can be challenging to accomplish all four goals at the same time, we do have some intuition on how to use the above framework to achieve any three of them. To meet the objective of (1), (2) and (4), while sacrificing effort elicitation, one can apply a flat-fee mechanism which compensates each agent for her cost of effort.¹ A similar approach is utilized in platforms like Amazon Mechanical Turk. We further use an example to illustrate how to use linear payment functions to elicit effort and truthful reporting, while sacrificing at least one of the other two goals. Suppose a desired equilibrium (e.g. all agents working with full effort) scores an agent 10 in expectation, a possible deviation (e.g. working with 90% of effort) scores her 9.9 in expectation, and (due to the variance) the minimum score is 0 in both cases. Suppose exerting full effort costs the agent \$10 worth of effort while exerting 90% of effort costs the agent \$9. In this example, to achieve (2), (3) and (4), the requester can first subtract a constant, i.e. 9.9, from every agent’s performance score, and scale it by 10. In this way, full effort can be elicited in equilibrium and every agent is paid \$10 in expectation, which is exactly the cost of effort. However, this violates limited liability. Instead, to achieve (1), (3) and (4), the requester has to directly scale the performance score by 10, which results in a payment of \$100 for each agent, ten times more than necessary. Such a problem is especially troublesome for performance measurements whose scores are unbounded below.²

Before we present our results, we first note that to put forth theoretical analysis, we assume that the noise of the performance measurement follows the Gaussian distribution whose mean and standard deviation are functions of agents’ effort.³ The Gaussian assumption is favored because it has only two parameters, yet they are both informative. Furthermore, our experiments using a real-data estimated agent-based model (ABM) suggest that the Gaussian model is a good fit for most commonly used performance measurements.

1.1 Our Results

As we have seen, linear payment functions preserve the truthfulness of performance measurements, but are not efficient in eliciting effort. In this paper, we propose using a non-linear rank-order (RO) payment function. Such a payment function is particularly useful in the peer prediction setting where an agent’s performance score depends on others’ reports, making it unfair to base the payments solely on the absolute values of the performance scores. Furthermore, RO-payment functions are easier to implement and trivially bound the ex-post budget.

To reveal the power of our approaches, this paper is structured as follows. In Section 4, we show how to optimize the RO-payment function to achieve (1), (2) and (3), and how to optimize a performance measurement under the Gaussian model. Then, in Section 5, we address the challenge of preserving truthfulness under the RO-payment function, which allows us achieve all four objectives simultaneously. Finally, in Section 7 and 8, we use agent-based model experiments to compare the performance of RO-payment function with the linear payment function, and evaluate several commonly used performance measurements. More details are shown below.

Optimizing the Payment Mechanism. We first assume that agents are truthfully reporting. As a running example, suppose a principal wants to recover the ground truth of a batch of tasks using the collected labels from a group of homogeneous agents, who have the same utility function and information structure.⁴ The principal’s problem is to design a payment mechanism to minimize the expected cost of budget for eliciting a goal effort in the symmetric equilibrium, i.e. no unilateral deviation in effort can increase an agent’s expected utility.⁵

¹Note that the flat-fee mechanism is not strictly truthful, as defined in the peer prediction literature [28, 22], in the sense that it does not ensure truth-telling to be a better-paid equilibrium than any uninformative reporting strategy profile.

²Because there does not exist an affine transformation to guarantee limited liability, (1) and (4) cannot be obtained at the same time.

³Due to the complexity of the performance measurements, it is theoretically challenging to analyze their performance score distribution, so some distribution-level assumptions are necessary.

⁴Although not without loss of generality, homogeneous agents are widely assumed in the principal-agent literature [9, 25]. The selection process, both self-selection and that executed by the principal (e.g. filtering on the platform) could result in increased homogeneity among agents’ background. Furthermore, agents are homogeneous while dealing with objective tasks with low dependence on experience.

⁵Symmetric equilibria are commonly used in economics literature [23, 8, 17] due to their tractability and analytical insights. In the settings we envision these being used, asymmetric equilibria are often closely approximated by symmetric equilibria. This is because agents can play a random strategy from an asymmetric equilibrium and, as the system grows, little changes.

First, given a performance measurement, we present analytical solutions for the optimal RO-payment functions for three types of agents: risk/loss-neutral, risk-averse, and loss-averse. Although similar problems have been studied in the economics literature as tournaments [8, 17], our results address a gap by incorporating individual rationality (IR) as a hard constraint in the optimization problem, which is an important condition in the crowdsourcing setting. Requiring the function to pay agents at least their cost of effort, we find that IR, while binding, results in optimal RO-payment functions that are more inclusive (rewarding more agents).

Second, given a RO-payment function, we examine how to optimize the performance measurement. Under the Gaussian model with an additional assumption that any unilateral deviation in effort only shifts the score distribution without changing its shape, we identify a sufficient statistic called the *sensitivity*. The sensitivity serves as a new criterion for evaluating a performance measurement: the higher the sensitivity, the lower the required payment for eliciting a desired effort. In general, a performance measurement with higher sensitivity is more accurate (has lower variance) and more sensitive to changes in effort.

Truthfulness Under the Rank-Order Payment Function. Although the optimized RO-payment function is effective in eliciting effort, it does not preserve the truthfulness of the performance measurement. As an example, under the winner-take-all tournament, an untruthful reporting strategy that reduces the expected performance score but increases its variance, can improve the chances of winning the top prize. Therefore, a truthful payment function must penalize the incentive to increase variance at the cost of decreased expected score. We propose adding a common noise to each agent’s performance score. Under the Gaussian model, we prove that adding a zero-mean Gaussian noise can help guarantee truthfulness in a winner-take-all tournament. However, the added noise decreases the sensitivity of a performance measurement. This observation suggests a new property of a performance measurement — the *variational robustness* — which quantifies how much noise is required to guarantee truthfulness of a performance measurement under the RO-payment function. Our agent-based model experiments suggest that most of the commonly used performance measurements have high variational robustness. Compared with the linear payment functions, we empirically show that the optimized RO-payment functions are significantly more effective in eliciting effort even after adding noise to guarantee truthfulness.

Evaluating Realistic Performance Measurements. In practice, we are curious about which performance measurement should be applied for rewarding agents. Our paper puts forward a new dimension of evaluation: the ability of a performance measurement to incentivize a desired level of effort at a low cost. We show that two properties matter: sensitivity, which measures how much the performance score changes with respect to the change of effort, and variational robustness, which captures the ability of a performance measurement preventing untruthful strategies from increasing the variance of the performance score. In this paper, we implement several state-of-the-art spot-checking and peer prediction mechanisms, and use real-world data estimated agent-based model to empirically evaluate them in terms of sensitivity and variational robustness. Our agent-based model results provide valuable insights into which mechanisms are most suitable for practical crowdsourcing settings.

2 Related Work

Tournament Design. The most relevant related works with respect to effort elicitation lie in the economics literature on tournaments. In line with our results, the winner-take-all (WTA) mechanism is proven to be optimal for neutral agents in small tournaments with symmetrically distributed noise [23], and in arbitrarily-sized tournament when the noise has increasing hazard rate [8]. The follow-up work [7] shows in the tournament setting that the equilibrium effort decreases as the noise of the effort measurement becomes more dispersed, in the sense of the dispersive order. In our situation with Gaussian noise, where the mean of the performance score is a function of effort, we show that it is the ratio of the derivative of the mean to the variance of the noise that affects the principal’s utility.

Green and Stokey [14] compare tournaments with independent contracts which pay agents based on their numerical outputs rather than the ranking of the outputs. In their model where the outputs of agents depend not only on their effort but also on an unknown common shock, they show that if there is no common

shock, the independent contracts dominant tournaments. However, if the distribution of the common shock is sufficiently diffuse, tournaments dominant independent contracts.

For risk-averse agents, Krishna and Morgan [23] show that the optimal RO-payment function is WTA when there are $n \leq 3$ risk-averse agents, and should pay the agent ranked in the second place positively when $n = 4$. Kalra and Shi [17] show that, for an arbitrary number of agents, the more risk-averse the agents are, the larger the number of agents should be rewarded with a focus on logistic and uniform noise distributions. Drugov and Ryvkin [6] generalize their results by considering more general noise distributions and non-separable preferences.

We note that in tournament design, the IR constraint is usually buried into the sufficient conditions for the existence of equilibrium. However, when considering IR, the optimal payment function remains unknown. This problem is essential in our crowdsourcing setting where IR is usually a binding constraint.

Forecast Competition. Forecast competition studies the problem of how to optimally select a winner from a pool of contestants based on prediction accuracy. Similar to a winner-take-all setting, the forecasters may have incentive of misreporting so as to increase the variance of its scores, even at a cost of the expected score. Witkowski et al. [32] propose the Event Lotteries Forecast (ELF) mechanism, which is shown to be approximately truthful when there are enough events. Frongillo et al. [11] further reduce the required number of events to $O(\log n)$ where n is the number of forecasters. The main idea of preserving truthfulness in their setting is to turn the selection of forecasters into a lottery, where the probability of winning is divided among all participants. Inspired by their solution, we show that adding noise to agents' performance scores can help preserve the truthfulness of the performance measurement.

Crowdsourcing and the Principal-Agent Problem. There exists more literature that considers the crowdsourcing problem from the principal-agent perspective. Ho et al. [16] model the crowdsourcing process as a multi-round principal-agent problem. Ghosh and Hummel [13] consider agents with heterogeneous ability and endogenous effort and focus on how these factors affect the optimal contract for the principal. The main difference between their work and ours is that they do not consider the payments to the agents as a cost of the principal (e.g. the payments are unredeamable points), unlike in our setting where crowd work is compensated with money. Easley and Ghosh [9] consider a crowdsourcing model where agents are strategic in deciding whether to participate in a task. Similar to Green and Stokey [14], they focus on when the principal should apply an output-independent contract or a winner-take-all tournament, which is shown to depend on the agents' behaviour models.

The principal-agent literature often assumes that the relationship between worker effort and outputs is known, however, this is usually not the case in practice. To address this issue, Kaynar and Siddiq [18] propose a non-parametric model for estimating effort-outcome distributions from principal-agent datasets. The model is proven to be statistically consistent. The experiments conducted using Amazon Mechanical Turk, in conjunction with the proposed estimation method, provide evidence for the effectiveness of performance-based incentive schemes in promoting effort.

Spot-Checking and Peer Prediction. Literature on spot-checking and peer prediction focuses on designing truthful mechanisms (e.g. whether agents can benefit by manipulating their reports) mostly in the binary-effort setting [12, 4, 28, 22]. Kong and Schoenebeck [20] consider a discrete hierarchical effort model where choosing higher effort is more informative but more costly. With assumptions, the maximum effort is proven to be elicitable and payments are optimized using a linear program that requires detailed knowledge of agent costs and quality. Recent works mostly study how to obtain stronger truthful guarantee with fewer samples [26, 19] and how to deal with different types of agents [1, 27].

Our approach diverges sharply from previous peer prediction work which focuses nearly entirely on strategic considerations where linear rescaling is the only known technique available. Instead, we separate the agent's choices of how much effort to exert from how honestly to report. This allows us to use a principal-agent framework to study how to elicit effort. In general, we obtain a weaker truthfulness guarantee, which is derived from empirical results showing that a litany of strategies do not work. Of course, it is possible that some strategy we failed to consider does work. However, our truthfulness results are slightly stronger in several ways as well. First, we need not rely on the number of agents going to infinity, but can run tests in finite settings (note that some mechanisms that are provably truthful in the limit of a large number of agents and conceivably

allow beneficial manipulations in the finite settings we consider). Additionally, some of the mechanisms do not work for all prior distributions, and we can empirically test if they work for the prior distributions that we learn from data. In general, we obtain a weaker truthfulness guarantee, where the theory is developed under the Gaussian assumption and the winner-take-all tournament, and the robustness of our idea largely depends on the empirical results. However, our truthfulness results are stronger in several ways as well. For instance, our method preserves truthfulness not only for linear payment functions, but also for non-linear RO-payment functions, and we believe the same idea can be extended to other payment functions like threshold functions. Additionally, some peer prediction mechanisms only have truthfulness guarantees under certain assumptions, such as a large or even infinite number of agents [22]. By testing truthfulness on synthetic data, we provide valuable insights into how well the mechanisms perform in practice.

3 Model

In Section 3.1, we introduce the basic concepts of the crowdsourcing problem and present a crowdsourcing model that is mainly used to generate synthetic data for our agent-based model experiments in Section 6. Next, in Section 3.2 and 3.3, we map the crowdsourcing problem into a principal-agent problem using the Gaussian assumption. We note that our theoretical analysis in Section 4 and 5 does not rely on the crowdsourcing model, in particular, the assumptions of signals and reports.

Throughout the paper, we use capital letters to denote random variables and lowercase letters to denote their specific values. We use bold font to denote vectors or matrices.

3.1 Crowdsourcing

A principal (requester) has a set of m tasks $[m] = \{1, 2, \dots, m\}$. Each task $j \in [m]$ has a ground truth $y_j \in \mathcal{Y}$ —that the principal would like to recover—which was sampled from a prior distribution $w \in \Delta_{\mathcal{Y}}$, where \mathcal{Y} is a discrete set and $\Delta_{\mathcal{Y}}$ is the set of all possible distributions over \mathcal{Y} . To this end, each task is assigned to n_0 agents and each agent i is assigned a subset of tasks $A_i \subseteq [m]$. Let m_a denote the maximum number of tasks assigned to any agent. Then $|A_i| \leq m_a$ for every agent i . This implies a lower bound on the number of agents: $n \geq \lceil m \cdot n_0 / m_a \rceil$. Let n be the number of agents.

Effort and cost. Agents are strategic in choosing an effort level. Let $e_i \in [0, 1]$ denote the effort chosen by agent i . Let $c(e)$ be a non-negative, increasing and convex cost function.

Signals and reports. Each agent i working on an assigned task j , receives a signal denoted $X_{i,j} \in \mathcal{X}$, where \mathcal{X} is the signal space. We assume that $0 \notin \mathcal{X}$ and let $X_{i,j} = 0$ for any $j \notin A_i$. For task $j \in A_i$, $X_{i,j}$ are i.i.d. sampled from a distribution that depends only the ground truth y_j and agent i 's effort level e_i .

Let Γ_{work} and Γ_{shirk} be $|\mathcal{Y}|$ by $|\mathcal{X}|$ matrices, where, for $y \in \mathcal{Y}$ and $s \in \mathcal{X}$, the y, s entry of Γ_{work} and Γ_{shirk} denotes the probability that an agent who puts in full effort and no effort, respectively, will receive a signal s when the ground truth is y . Given e_i , agent i 's signal $X_{i,j}$ for the j th task where the ground truth is y_j will be sampled according to the y_j th row of

$$e_i \Gamma_{\text{work}} + (1 - e_i) \Gamma_{\text{shirk}}.$$

We will let Γ_{shirk} be uniform in each column. This setup is a modified version of the Dawid-Skene (DS) model [5] where we have added effort.

We use \mathbf{x} and $\hat{\mathbf{x}}$ to denote the signal and report profiles of all agents respectively. Note that $\hat{\mathbf{x}}$ is not necessarily equal to \mathbf{x} when agents are strategically reporting. Unless otherwise, we assume all agents report truthfully, so that $\hat{\mathbf{x}} = \mathbf{x}$. Strategic reporting is discussed in Section 3.3.

Mechanism. Given $\hat{\mathbf{x}}$, a payment mechanism $\mathcal{M} : (\{0\} \cup \mathcal{X})^{n \times m} \rightarrow \mathbb{R}_{\geq 0}^n$ pays each agent i a non-negative payment t_i . We decompose the payment mechanism into two parts (Fig. 1). First, we apply a performance measurement $\psi : (\{0\} \cup \mathcal{X})^{n \times m} \rightarrow \mathbb{R}^n$ on agents' reports that outputs a (possibly negative and random) score $s_i = \psi(\hat{\mathbf{x}})_i$ for each i . In our experiments, we focus on two sets of performance measurements: spot-checking and peer prediction, which will be discussed later.

Second, we apply a rank-order payment function that pays \hat{t}_j to the j 'th ranked agent according to her performance score. WLOG, suppose $s_1 \geq s_2 \geq \dots \geq s_n$. Then, agent i 's payment is $t_i = \hat{t}_i$. As a comparison, in Section 7.2, we consider a linear payment function as a baseline that rewards each agent i a linear transformation of her performance score, i.e. $t_i = a \cdot s_i + b$ where a and b are constants.

Definition 3.1. We call a RO-payment function increasing if $\hat{t}_j \geq \hat{t}_k$ if $j \leq k$.

3.2 The Principal-Agent Model

We seek a payment mechanism that maximizes the principal’s payoff in the symmetric equilibrium. Now, assuming agents report their signals truthfully, we model this crowdsourcing problem as a principal-agent problem.

First, the principal assigns the tasks to agents and commits to a payment mechanism consisting of a performance measurement ψ and a RO-payment function \hat{t} . Then, the agents respond by choosing an effort level that maximizes their expected utility. Intuitively, the effort affects the distribution of an agent’s performance score, which in turn affects the rankings and payments. In this paper, we consider three utility functions: In this paper, we examine the robustness of our results by considering three utility functions, in order to take into account realistic considerations of the risk and loss preferences of agents in the real-world:

$$u_a(t_i, e_i) = \begin{cases} t_i - c(e_i) & \text{for neutral agents,} \\ t_i - c(e_i) - \rho \cdot (c(e_i) - t_i)^+ & \text{for loss-averse agents,} \\ r_a(t_i) - c(e_i) & \text{for risk-averse agents.} \end{cases}$$

Here, for loss-averse agents⁶, $(x)^+$ equals x for $x \geq 0$ and 0 otherwise, and ρ is a non-negative loss-aversion factor. For risk-averse agents, r_a is a non-negative, concave and differentiable function with $r_a(0) = 0$ and $r'_a(0) < \infty$. Loss-averse agents incur an additional loss when they are not compensated for the work they expend. Risk-averse agents proportionally value moderate rewards more the high rewards.

For simplicity, we often use the following utility function which contains all three types of utility functions as a special case,

$$u_a(t_i, e_i) = r_a(t_i) - c(e_i) - \rho \cdot (c(e_i) - t_i)^+. \quad (1)$$

We focus on the solution concept of symmetric equilibrium: all agents exerting effort ξ is an equilibrium if any unilateral deviation by an agent will decrease their expected utility, i.e. $\mathbb{E}[u_a(t_i(e_i, \xi), e_i)] \leq \mathbb{E}[u_a(t_i(\xi, \xi), \xi)]$ for any $e_i \in [0, 1]$, where $t_i(e_i, \xi)$ is a random payment function on agent i ’s effort and all the other agents’ effort $e_k = \xi$ for any $k \neq i$.

The problem of the principal is then to optimize the payment mechanism such that a goal effort ξ can be incentivized in the symmetric equilibrium with the minimum payment.⁷ Additionally requiring the payment to satisfy limited liability (LL) and individual rationality (IR) leads us to the principal’s optimization problem. In Appendix C, we further provide a variant of the principal’s model, where the principal also cares about the fairness of the rank-order payments, e.g. reducing the variance of the payments.

Remark (Individual rationality). To get a sense for the IR constraint, consider the following situation: The principal would like 10 agents to each exert \$10 of effort. Under a winner-take-all contest, the principal rewards the top agent \$80 which induces a symmetric equilibrium where the agents each contribute \$10. This is not IR because each agent gets a utility of $-\$2$. However, simply increasing the payment of the top agent changes the effort in equilibrium. A different payment structure is needed to achieve IR with the target effort.

The Gaussian assumptions. However, the optimization problem over the space of all performance measurements is still too hard to analyze.⁸ To make it theoretically tractable, as commonly assumed in principal-agent literature, we apply the Gaussian noise assumption. Again, let e_i be agent i ’s effort and ξ be all the other agents’ effort.

⁶Note that we mainly consider the 1-order loss-aversion. We briefly discuss the case of higher order loss-aversion, i.e. $u_a(t_i, e_i) = t_i - c(e_i) - \rho \cdot ((c(e_i) - t_i)^+)^r$ for $r > 1$, in Section 4.1.2.

⁷We note that in reality, the principal can optimize over the space of the parameters of the crowdsourcing system such as the number of agents, the number of tasks each agent answers and the goal effort. However, the optimization over these parameters requires a finer-grained model of the principal’s utility, i.e. how does the principal evaluate the contributions from agents, which is beyond the interests of this paper. Therefore, we assume the principal fixes these dimensions and knows the goal effort that he wants to elicit. We will show, later in this paper, how the goal effort affects the principal’s decision.

⁸The main difficulty is that, in general, we don’t analytically know what’s the distribution of the performance scores output by a performance measurement, which usually has no closed-form.

Assumption 3.1. We assume the agent i 's performance score S_i follows the Gaussian distribution with p.d.f. $g_{e_i, \xi}^{(i)}$ and c.d.f. $G_{e_i, \xi}^{(i)}$, where the mean $\mu(e_i, \xi)$ and standard deviation $\sigma(e_i, \xi)$ are functions of agents' effort. Furthermore, let $g_{e_i, \xi}^{(-i)}$ and $G_{e_i, \xi}^{(-i)}$ be the same notations for all the other agents' score distribution under the same effort profile. We assume μ and σ to be differentiable.

Assumption 3.2. We assume the distribution $g_{e_i, \xi}^{(-i)}$ is independent of e_i .

Assumption 3.2 implies that any unilateral deviation $e_i \in [0, 1]$ from a symmetric effort profile where all agents' effort is ξ will not change other agents' score distribution. This implies $g_{e_i, \xi}^{(-i)} = g_{\xi, \xi}^{(-i)}$. This assumption intuitively holds for spot-checking mechanisms, where agents' performance scores are independent conditioned on the ground truth, and for peer prediction mechanisms when the number of agents is large (in which case, the influence of any one agent's effort is diluted by the number of agents). Our empirical examination shows that Assumption 3.2 holds for a reasonable number of agents in realistic crowdsourcing settings, e.g. $n = 50$. For simplicity, while fixing ξ , we use g_{e_i} to denote agent i 's score distribution and g_ξ to denote other agents' score distribution.

We additionally make the following assumption which, at a high level, guarantees that a unilateral deviation to a higher effort does not harm the agent's expected performance score.⁹

Assumption 3.3. Fixing ξ , let $\mu'_\xi(e_i) = \frac{\partial \mu(e_i, \xi)}{\partial e_i}$ be the derivative of $\mu(e_i, \xi)$ over e_i as a function of e_i and $\sigma'_\xi(e_i)$ is the similar notation for the standard deviation. We assume $\mu'_\xi(e_i) \geq 0$ and furthermore, $\mu'_\xi(e_i) + \sigma'_\xi(e_i) \geq 0$ for any $e_i, \xi \in [0, 1]$.

The above assumption is necessary for establishing our theoretical results. Intuitively, as our ABM experiments demonstrate, the derivative of the standard deviation is of lower order than the derivative of the mean. Therefore, Assumption 3.3 is approximately saying that μ is increasing in e_i .

3.3 Strategic Reporting

We consider strategic reporting in Section 5 and define the terminologies here. While considering strategic reporting, we fix all agents' effort and omit the notations that indicate agents' effort. Then, for agent i , given her signal $X_{i,j}$ on task j , let $\hat{X}_{i,j} \in \mathcal{X}$ be her report on that task. As a common assumption in peer prediction [4, 1], we assume task-independent strategies, which imply that agent i will first choose a reporting strategy $\theta_i : \mathcal{X} \rightarrow \Delta_{\mathcal{X}}$, then draw $\hat{X}_{i,j}$ from the distribution $\theta_i(X_{i,j})$ as her report for every assigned j .¹⁰ In other words, fixing a reporting strategy, the report distribution of one task depends only on the ground truth of that task per se. We assume agents have a common strategy space Θ that is compact. Specifically, we use τ_i to denote the truth-telling strategy, i.e. $\tau_i(X) = X$.

Previous literature has provided a large number of choices of truthful performance measurements. We first note that the truthfulness of a performance measurement is defined on the expected scores, i.e.

$$\mathbb{E}[S_i(\theta_i, \theta_{-i})] = \mathbb{E}_{\mathbf{X}, \theta} [\psi(\hat{\mathbf{X}})],$$

where ψ is the performance measurement.

Definition 3.2. We call a performance measurement truthful if no unilateral deviation from the truth-telling strategy profile can increase the expected performance score, i.e. $\mathbb{E}[S_i(\theta_i, \tau_{-i})] \leq \mathbb{E}[S_i(\tau_i, \tau_{-i})]$ for any agent i and strategy $\theta_i \in \Theta$. Furthermore, a performance measurement is strongly truthful if the above inequality is strict.

However, what we want is the truthfulness of a payment mechanism, which should guarantee equilibrium in terms of the payments (not scores). For a given performance measurement, let $p_j(\theta_i, \theta_{-i})$ be the probability that the performance score of agent i is ranked in the j 'th position. The expected payment under a strategy profile θ can be written as

$$\mathbb{E}[t_i(\theta_i, \theta_{-i})] = \sum_{j=1}^n p_j(\theta_i, \theta_{-i}) \hat{t}_j.$$

⁹In our experiments, we observe that $\sigma'_\xi(e_i)$ is insignificant compared with $\mu'_\xi(e_i)$.

¹⁰A recent work [33] generalizes the design of peer prediction mechanisms to task-dependent strategies.

Definition 3.3. We call a payment mechanism truthful if no unilateral deviation from the truth-telling strategy profile can increase the expected payment, i.e. $\mathbb{E}[t_i(\theta_i, \tau_{-i})] \leq \mathbb{E}[t_i(\tau_i, \tau_{-i})]$, for any agent i and strategy $\theta_i \in \Theta$. Furthermore, a payment mechanism is strongly truthful if the above inequality is strict.

Note that the linear payment function trivially transfers the truthfulness of the performance measurement to the truthfulness of the payment mechanism. However, because the rank-order payment function is non-linear, this property does not generally hold. We will study this issue in depth in Section 5.

The Gaussian assumptions. Again, to theoretically track the problem, we adopt the following two assumptions which are analogous to Assumption 3.1 and 3.2 but with respect to agents' reporting strategies.

Assumption 3.4. We assume the agent i 's performance score S_i follows the Gaussian distribution with p.d.f. $g_{\theta}^{(i)}$ and c.d.f. $G_{\theta}^{(i)}$, where the mean $\mu_i(\theta)$ and the standard deviation $\sigma_i(\theta)$ are functions of agents' strategy profile. Furthermore, we assume the domains of μ_i and σ_i are compact for any i .¹¹

Assumption 3.5. Let θ be the initial strategy profile. Suppose agent i unilaterally deviates to an arbitrary strategy θ'_i . Let the corresponding change in the mean of the score distribution of an agent $j \in [n]$ be $\Delta\mu_j(\theta'_i, \theta) = \mu_j(\theta'_i, \theta_{-i}) - \mu_j(\theta)$. We assume $|\Delta\mu_i(\theta'_i, \theta)| \geq \Delta\mu_j(\theta'_i, \theta)$ for any $j \neq i$ and $\theta'_i \in \Theta$.

Assumption 3.5 indicates that if an agent unilaterally changes her reporting strategy, she will change the mean of her own score more than the mean of any other agent's score. Note that this assumption is weaker than Assumption 3.2 - the analogue assumption for agents' effort strategy - as the latter requires all the other agents' score distributions stay unchanged if there is an unilateral deviation in effort. Again, we find that this assumption mild which generally holds for about 50 agents.

Remark. We note that our theoretical results are developed in the "idealized setting" where Assumption 3.1 - 3.5 hold. In particular, we explicitly assume that Assumption 3.1 - 3.3 hold in Section 4, and Assumption 3.4 - 3.5 hold in Section 5. We further call a performance measurement that satisfy these assumptions an idealized performance measurement.

4 Optimizing the Payment Mechanism

This section focuses on designing payment mechanisms with the objectives of limited liability (1), budget efficiency (2) and (3) effort elicitation, while we ignore the truthful guarantee (4) of the performance measurement. In particular, we show how to reward agents optimally for a desired effort level in the idealized setting. The optimization consists of two parts: optimizing the rank-order payment function while fixing any idealized performance measurement, and optimizing the idealized performance measurement given the optimal RO-payment function. For the former, we observe that the optimal RO-payment function is increasing for all agent utility models that we considered, and both risk/loss-aversion and individual rationality will make the optimal RO-payment function more inclusive which rewards a larger number of agents. For the latter, we identify the sufficient statistic of a good performance measurement called the sensitivity. We show that a performance measurement with higher sensitivity can incentivize the same effort level in the symmetric equilibrium with a lower total payment. Our discussions in this section provide suitable solutions for payment mechanism design in cases where strategic manipulations are not a significant issue, while we defer the discussion of additional guaranteeing truthfulness to Section 5.

4.1 Optimizing the Rank-Order Payment Function

We first rewrite the principal's problem given a performance measurement ψ . Suppose all the agents except i exert an effort ξ . Then, given ψ , agent i knows the probability that she will end up with each rank j when her effort is e_i , which is denoted as $p_j(e_i, \xi)$. Recall that G_{e_i} and G_{ξ} are the c.d.f. of the score distribution of agent i and all the other agents respectively. Then,

$$p_j(e_i, \xi) = \binom{n-1}{j-1} \int_{-\infty}^{\infty} G_{\xi}(x)^{n-j} [1 - G_{\xi}(x)]^{j-1} dG_{e_i}(x). \quad (2)$$

¹¹Given the compact strategy space and the finite signal space, this is a mild assumption.

We then can write agent i 's expected utility under the RO-payment function \hat{t} as

$$\mathbb{E}[U_a(e_i, \xi)] = \sum_{j=1}^n p_j(e_i, \xi) u_a(\hat{t}_j, e_i), \quad (3)$$

where U_a denotes the random variable of the agent's utility and u_a is agent's utility function defined in Eq. (1).

Maximizing the expected utility w.r.t. e_i then leads to the first order constraint (FOC) which is a necessary condition of symmetric equilibrium. For sufficiency, additional conditions on the distribution of the performance score and the agents' cost function are required. For example, it is shown that when the distribution of the noise (in our case, this is the Gaussian) is "dispersed enough", the existence of symmetric equilibrium is guaranteed [24]. In our theory sections, we assume FOC is sufficient for symmetric equilibrium, which implies that the local maximum of the agent's utility function is also the global maximum. Our empirical results further verify that this is indeed the case for the considered performance measurements and cost functions.

Let $p'_j(\xi) = \frac{\partial p_j(e_i, \xi)}{\partial e_i} \big|_{e_i=\xi}$ denote the derivative of the probability an agent ends up with rank j w.r.t. a unilateral deviation in effort when all agents' effort is ξ , and let $c'(\xi)$ denote the derivative of the cost. Also, note that $p_j(\xi, \xi) = \frac{1}{n}$ for any j due to symmetry. Now, given n and ξ , we formally write down the principal's problem.

$$\begin{aligned} \min_{\hat{t}} \quad & \sum_{j=1}^n \hat{t}_j \\ \text{s.t.} \quad & \hat{t} \geq 0 \quad (LL), \quad \frac{1}{n} \sum_{j=1}^n u_a(\hat{t}_j, \xi) \geq 0 \quad (IR), \quad \sum_{j=1}^n p'_j(\xi) u_a(\hat{t}_j, \xi) = 0 \quad (FOC). \end{aligned} \quad (4)$$

Before we present our results, we present the following lemma that is essential for future proofs.

Definition 4.1. The derivative of the probability of ranking, \mathbf{p}' , is said to have *rank-order impact* at ξ , if $p'_j(\xi)$ is decreasing in j for any $1 \leq j \leq n$.

Lemma 4.2. Fixing $\xi \in [0, 1]$, if $n \rightarrow \infty$, \mathbf{p}' has rank-order impact at ξ .

We leave the proof in Appendix A.1. Lemma 4.2 shows that after convergence, a small unilateral deviation results in a probability of ranking that is monotone decreasing in j . The key of the proof lies in the fact that after convergence, $p'_j(\xi)$ can be approximated with some form of the quantile function of Gaussian, which is known to be the inverse error function. Then, with the monotonicity of the inverse error function, we complete the proof.

Next, we present the optimal RO-payment functions for the principal's problem under the three utility models.

4.1.1 Neutral Agents

Now suppose agents are neutral, i.e. $u_a(t_i, e_i) = t_i - c(e_i)$. We have the following results.

Proposition 4.3. Suppose \mathbf{p}' has rank-order impact at $\xi \in [0, 1]$, and agents are neutral.

1. **IR is not binding:** If $\frac{c'(\xi)}{p'_1(\xi)} \geq n \cdot c(\xi)$, the optimal RO-payment function is winner-take-all, i.e. $\hat{t}_1 = \frac{c'(\xi)}{p'_1(\xi)}$ is the reward to the top one agent and $\hat{t}_j = 0$ for $1 < j \leq n$;
2. **IR is binding:** Otherwise, the optimal RO-payment function is not unique and can be achieved by a threshold function that rewards the top \hat{n} agents equally, i.e. $\hat{t}_j = \frac{n}{\hat{n}} c(\xi)$ for $1 \leq j \leq \hat{n}$ and 0 otherwise. The threshold \hat{n} is determined by $\frac{n}{\hat{n}} \sum_{j=1}^{\hat{n}} p'_j(\xi) c(\xi) = c'(\xi)$.

The proof is deferred to Appendix A.2. As a sketch, the proposition holds because by Theorem 4.2 when \mathbf{p}' has rank-order impact, $p'_j(\xi)$ is decreasing in j . This implies that if IR is not binding, when we take the gradient of the total payment in Eq. (4) w.r.t. each \hat{t}_j , the gradient reaches its maximum when $j = 1$. Thus, the most payment-saving RO-payment function is to put all of the budget on \hat{t}_1 to maximize the gain of any unilateral deviation to a higher effort.

It is worth noting that unless $\frac{c'(\xi)}{p'_1(\xi)} \rightarrow \infty$, Theorem 4.3 implies that IR is always binding. However, we emphasize that the condition $n \rightarrow \infty$ in Theorem 4.3 (and subsequent propositions) is only required by the proof of Lemma 4.2. Empirical results suggest that the lemma still holds for a reasonable group size, e.g. $n = 50$ (as outlined in Section 6.2). This implies that the winner-take-all payment function can be optimal (corresponding the case that IR is not binding) for a larger range of settings. Combining Lemma 4.2 and Proposition 4.3 immediately gives us a corollary suggesting that when $n \rightarrow \infty$, the optimal RO-payment function for neutral agents follows the structure presented in Proposition 4.3. We note that although $n \rightarrow \infty$ is sufficient to prove the rank-order impact of \mathbf{p}' , it is not necessary. Our empirical results suggest that the rank-order impact still holds for a reasonable group size, e.g. $n = 50$ (as outlined in Section 6.2).

We further note that when agents are neutral and IR is binding, the minimum total payment equals the total cost $n \cdot c(\xi)$. We name the payment that makes IR binding the *IR-minimal payment*. When agents are neutral, the IR-minimal payment can be achieved with multiple payment functions, where we show that threshold functions can be one of the solutions.

4.1.2 Loss-averse and Risk-averse Agents

Suppose agents are loss-averse, i.e. $u_a(t_i, e_i) = t_i - c(e_i) - \rho \cdot (c(e_i) - t_i)^+$. The analytical solution of the optimal RO-payment function becomes more complicated in this case. Here, we present a simplified version of our result while leaving the detailed version in Appendix A.3.

Proposition 4.4. (Simplified) Suppose \mathbf{p}' has rank-order impact at $\xi \in [0, 1]$, and agents are loss-averse.

1. **IR is not binding:** The optimal RO-payment function pays 1) 0 to the bottom agents with ranking $j > \bar{n}$, 2) $c(\xi)$ to the intermediate agents with ranking $1 < j \leq \bar{n}$, and 3) $\hat{t}_1 > c(\xi)$ to the top one agent. Here, the threshold $\bar{n} \leq \frac{n}{2}$ is determined by $(1 + \rho)p'_n(\xi) = p'_1(\xi)$;
2. **IR is binding:** The optimal RO-payment function follows the same structure as the case of IR not binding, but with a threshold $\hat{n} \geq \bar{n}$.

The proof is shown in Appendix A.3. As a sketch, note that the gradient of the total payment w.r.t. \hat{t}_j is maximized at $j = 1$ only when $\hat{t}_1 \leq c(\xi)$. When $\hat{t}_1 > c(\xi)$ the gradient is discounted with a factor $\frac{1}{1+\rho}$. Therefore, because $p'_j(\xi)$ is decreasing in j , the optimal RO-payment function will “fill in” \hat{t}_j to $c(\xi)$ in the increasing order of j until j is greater than \bar{n} in which case the discounted gradient w.r.t. \hat{t}_1 is larger the undiscounted gradient w.r.t. $\hat{t}_{\bar{n}}$. Then, the rest budget is put on \hat{t}_1 .

Proposition 4.4 shows that for loss-averse agents, the optimal RO-payment function has three levels of payments: the bottom agents are paid zero; intermediate agents receive the baseline payment that equals to their cost; the top one agent gets a bonus that is larger than her cost. We call this type of RO-payment function the *winner-take-more* payment function. Perhaps interestingly, winner-take-more takes a similar form of the baseline-bonus payment scheme which tends to perform well in real-world scenarios [15]. We show the optimal RO-payment functions in Fig. 2 to better illustrate our ideas.

To better illustrate our results, we introduce the inclusiveness of a (monotone) RO-payment function.

Definition 4.5. Given a monotone RO-payment function such that $\hat{t}_j \geq \hat{t}_k$ whenever $j \leq k$, the *inclusiveness* of such a RO-payment function is defined as the number of agents who receive non-zero payments, denoted as n^I . We call a RO-payment function more inclusive than another RO-payment function if the inclusiveness of the former is weakly larger.

For example, $n^I = 1$ for the winner-take-all tournament, and $n^I = \bar{n}$ and $n^I = \hat{n}$ in the case of loss-averse agents with IR not binding and binding respectively. Now, we show that n^I increases for the optimal RO-payment function as agents become more loss-averse.

Corollary 4.6. Suppose $n \rightarrow \infty$ and agents are loss-averse. The inclusiveness of the optimal RO-payment function n^I is (weakly) increasing in ρ .

Remark. Our results for loss-averse agents rely on the first-order loss-aversion model. The study of higher-order loss-aversion is beyond the scope of this paper, i.e. $u_a(t_i, e_i) = t_i - c(e_i) - \rho \cdot ((c(e_i) - t_i)^+)^r$ for $r > 1$. Our conjecture is that in this case, the optimal RO-payment function will no longer pay a fraction of agents constantly, but pay agents decreasingly w.r.t. their ranking.

The analogue propositions for risk-averse agents are left in Appendix A.4, while we use Fig. 2 to visually summarize the key insights of our results.

As shown in Fig. 2, Theorem 4.4 illustrates that the optimal RO-payment function for risk-averse agents also only rewards the agents whose ranking is above some threshold, and the rewards are decreasing in the ranking j . As a proof sketch, this is because if the principal keep increasing the payment for the top one agent, the gradient of the total payment w.r.t. \hat{t}_1 is discounted greater and greater due to the concavity of the reward function. Therefore, after \hat{t}_1 becomes large enough, the principal is better off to increase \hat{t}_2 , \hat{t}_3 and so on from zero to something instead of keep increasing \hat{t}_1 .

The proposition shows that, in line with the results of neutral and loss-averse agents, the payment function becomes more inclusive when IR is binding compared with the case where IR is ignored. First, similar to the results for loss-averse agents, the optimal payment for risk-averse agents also rewards the bottom agents zero; while the agents ranked above a threshold are rewarded strictly increasingly in terms of their rankings. Intuitively, for both types of agents, it becomes more efficient to reward high-effort behaviors by rewarding some lower-ranked agents than giving all the rewards to the top agent. Second, it is a common pattern for all types of agents that the payment function tends to be more inclusive when IR is binding compared to when it is ignored. This is because the optimal payment function in the case of IR binding must compensate agents' cost of effort. Compared with the optimal payment function when IR is ignored, a higher total payment is required. However, increasing the payments for high-ranked agents motivates the agents to deviate to a higher effort level. Thus, some lower-ranked agents who previously received zero payment are now rewarded positively.

It's worth noting that the result in Corollary 4.6 does not extend to risk-averse agents. In other words, having a higher level of risk-aversion does not always result in a more inclusive RO-payment function.¹² We provide a formal illustration of this in Appendix A.4.1. The reason for this failure to generalize is because the model of risk-averse agents has more flexibility, meaning that $r_a(x)$ can be "more risk-averse" in many different ways.




Figure 2: The optimal RO-payment functions for three types of agents. Blue payments are the cases where IR is ignored, and orange payments are the cases where IR is considered. All payment schemes are increasing.

To sum up, we show that the optimal RO-payment functions for three common models of agent utility are all increasing payment schemes. Furthermore, inclusive payments are likely to be referred for two reasons. On one hand, IR requires the optimal RO-payment function to be more inclusive (compared with the case where IR is ignored) so as to guarantee the existence of equilibrium. On the other hand, when agents have “fairness” concerns, e.g. they are risk or loss-averse, more inclusive RO-payment functions are optimal (compared with the neutral case, where WTA is optimal). In Section 7.1, we empirically show how the inclusiveness of the optimal RO-payment functions interacts with the cost function and equilibrium effort.

4.2 Optimizing the Performance Measurement

A performance measurement can affect principal's optimal utility by affecting $p_j(e_i, \xi)$. In the idealized setting, assuming all the other agents but i exert effort ξ , every performance measurement maps agent i 's effort e_i to a Gaussian distribution of her performance score with mean and std functions of e_i . We denote these two functions as $\mu(e_i, \xi)$ and $\sigma(e_i, \xi)$ respectively, which determine how good a performance measurement is in

¹²Actually, we show that “more risk-averse” leads to more inclusive optimal RO-payment function when IR is not binding, but both it is not true when IR is binding.

our setting. Let $\mu'_\xi(\xi)$ be the derivative of μ w.r.t. e_i when $e_i = \xi$. For any RO-payment function that is increasing (Definition 3.1), we identify the following sufficient statistic of the performance of a performance measurement, called the *sensitivity*. Note that by Assumption 3.2, the score distribution of any other people follows $\mathcal{N}(\mu(\xi, \xi), \sigma(\xi, \xi))$. The main result of this section is that under an additional assumption on the score distribution, we identify a sufficient statistics of the ability of eliciting effort at a low cost of a performance measurement. Let $\mu'_\xi(\xi)$ and $\sigma'_\xi(\xi)$ be the derivative of μ and σ w.r.t. e_i when $e_i = \xi$.

Assumption 4.1. *We assume $\mu'_\xi(e_i) \gg \sigma'_\xi(e_i)$ for any $e_i, \xi \in [0, 1]$.*

The above assumption says that if an agent deviates to a slightly higher effort, the change of the standard deviation of her performance score is negligible compared with the change of the mean. In other words, fixing all the other agents' effort, varying the effort of one agent basically only shifts the score distribution of that agent but does not change the shape of the score distribution. The intuition is that while increasing one agent's effort, the change of the absolute positions of performance scores (the first-order statistic) is more significant than the change of the relative positions (the second-order statistic). Although the assumption is non-trivial, we empirically find that the ratio between $\sigma'_\xi(e_i)$ and $\mu'_\xi(e_i)$ is usually less than 0.1. Now, we introduce the sufficient statistics, called the *sensitivity*.

Definition 4.7. The sensitivity of a performance measurement whose performance score distribution has a mean of $\mu(e_i, \xi)$ and a standard deviation of $\sigma(e_i, \xi)$ is defined as $\delta(\xi) = \frac{\mu'_\xi(\xi)}{\sigma'_\xi(\xi)}$.

The sensitivity is defined under the concept of symmetric equilibrium and is dependent on the effort in the symmetric equilibrium. At a high level, a performance measurement is more sensitive if it can generate scores that are more sensitive in effort change (μ' is large) and have high accuracy (σ is small). Also note that $\delta(\xi) \geq 0$ by Assumption 3.3 and 4.1.

Proposition 4.8. *Under Assumption 4.1, for any performance measurement and increasing RO-payment function, the minimum total payment $\sum_{j=1}^n \hat{t}_j$ is (weakly) decreasing in $\delta(\xi)$ while fixing any $\xi \in [0, 1]$.*

The proof is shown in Appendix A.5. At a high level, the intuition is that if an agent slightly increases her effort, it becomes easier for her to be ranked in higher places. This effect is amplified by a performance measurement with higher sensitivity. Therefore, with a more sensitive performance measurement, the first order constraint in Eq. (4) can be satisfied with lower payment. Because both of the other constraints are independent of performance measurements, this completes the proof.

Now, we have optimized the performance measurement and the RO-payment function separately. The following corollary fits the optimization results together.

Corollary 4.9. *Fixing a goal effort, let ψ' be a performance measurement with higher sensitivity than ψ . Let \hat{t}' and \hat{t} be their corresponding optimal RO-payment functions, respectively. Then the payment mechanism consisting of ψ' and \hat{t}' has lower minimal total payment than the payment mechanism consisting of ψ and \hat{t} in the symmetric equilibrium.*

The proof is straightforward by comparing three payment mechanisms: mechanism 1 is consisted of ψ' and \hat{t}' , mechanism 2 is consisted of ψ' and \hat{t} and mechanism 3 is consisted of ψ and \hat{t} . First, by our results in Section 4.1, both \hat{t}' and \hat{t} are increasing. Then, by Theorem 4.8, mechanism 2 should be cheaper to implement than mechanism 3. Furthermore, mechanism 1 must be cheaper than mechanism 2 because \hat{t}' is the optimal RO-payment function for ψ' , which completes the proof.

Sensitivity and The Optimal Payment. Here, we discuss how the sensitivity of a performance measurement affects the payment under the optimal RO-payment function. We consider neutral agents as an example. Fixing a goal effort ξ , if the sensitivity is low, the derivative of the probability of being ranked first, denoted as $p'_1(\xi)$, will be small. In this case, the condition of IR is not binding in Proposition 4.3 always holds, which means agents are paid more than their cost of effort in expectation. As the sensitivity increases, $p'_1(\xi)$ increases, resulting in a decrease in the required payment in the winner-take-all tournament, i.e. $\hat{t}_1 = \frac{c'_1(\xi)}{p'_1(\xi)}$. If the sensitivity keeps increasing, at some point, the payment to guarantee the equilibrium of effort ξ becomes insufficient to guarantee IR, which makes IR binding. Increasing the sensitivity beyond this point still reduces $p'_1(\xi)$, but it does not reduce the total payment any further, where we reach the IR-minimal payment. Furthermore, since $p'_1(\xi)$ decreases as the sensitivity increases, the optimal RO-payment function has to be more inclusive which rewards more than one agent to guarantee the effort equilibrium.

5 Truthful Winner-Take-All Tournaments

So far, we have shown how to optimize a payment mechanism to incentivize a desired effort level. In this section, we further investigate the problem of how to preserve the truthfulness of a performance measurement under the non-linear RO-payment function. In particular, we focus on the winner-take-all tournament and assume the effort (and thus the cost) of agents is fixed. In the idealized setting, we show that adding a large noise to agents' performance scores can discourage untruthful deviations even when the non-linear payment function is applied. We note that the analysis in this section generally holds for all three types of agents.

5.1 High Variance Benefits Deviations

To begin, it is important to understand why a truthful performance measurement does not imply a truthful payment mechanism under RO-payment functions. Although a truthful performance measurement guarantees that any deviation from the truth-telling strategy profile will decrease the expected performance score, under the non-linear RO-payment function, the expected payment may increase. For example, an untruthful strategy that decreases the expected score, but increases the variance of the performance score, can potentially help the agent to secure the top rank in a winner-take-all tournament. This observation is formalized in the following lemma.

Lemma 5.1. *Under the winner-take-all tournament, let $\tilde{t}(\mu_i, \sigma_i)$ be the expected payment of agent i when her score follows $\mathcal{N}(\mu_i, \sigma_i)$ and any other agent's score follows $\mathcal{N}(\mu, \sigma)$. If $\mu_i \leq \mu$, then*

1. \tilde{t} is increasing in μ_i and σ_i .
2. If $n \geq 3$, for any μ_i , there exists a σ_i such that $\tilde{t}(\mu_i, \sigma_i) > \tilde{t}(\mu, \sigma)$.

The proof follows by showing that the first derivatives of \tilde{t} with respect to μ_i and σ_i are positive, and showing that an agent can be ranked first with approximately 1/2 probability if her score has a large enough variance. Details are shown in Appendix B.1. Lemma 5.1 shows that while fixing the mean score of a deviation, increasing the variance of the performance score benefits the deviation, and eventually, such a deviation will outperform truth-telling. Furthermore, based on the lemma, we can immediately obtain the following proposition, which suggests that any untruthful strategy that does not increase the variance of the performance score will never outperform truth-telling.

Proposition 5.2. *Given a truthful performance measurement, for any unilateral untruthful deviation, if it (weakly) decreases the variance of the performance score, it will never make (all three types of) the deviating agent better-off. Under a truthful performance measurement and the winner-take-all tournament, if a unilateral untruthful deviation (weakly) decreases the variance of the performance score, it (weakly) decreases the expected payment.*

Proof. We know that under a truthful performance measurement, any untruthful deviation will decrease the expected performance score. Then, by Lemma 5.1, the smaller the variance, the smaller the probability of winning the first prize, and thus the smaller the utility of such a deviation. The proposition follows because even when an untruthful strategy has the same variance as truth-telling, it has a smaller expected score, and thus leads to a smaller probability of winning the first prize. \square

5.2 Adding Noise Helps Truthfulness

We propose a solution to the above problem. To guarantee truthfulness, by Proposition 5.2, we only have to deal with the untruthful strategies that increase the variance of the performance score. The key is to reduce the difference between the variance of the score distribution of truth-telling and that of an untruthful strategy. We propose a method of adding a common noise to every agent's performance score and then apply the rank-order payment function. We wrap this idea into a modified payment mechanism called the *manipulation-robust payment mechanism*.

First, a *manipulation-robust performance measurement* is constructed based on a truthful performance measurement with an additional step. After applying the initial performance measurement, which turns agents' reports into the performance scores, for each agent $i \in [n]$, we i.i.d. sample a score ϵ_i from a Gaussian

distribution $g_\epsilon = \mathcal{N}(0, \sigma_\epsilon)$. The new performance score of agent i is then $s'_i = s_i + \epsilon_i$. In other words, a manipulation-robust performance measurement outputs performance scores with additive Gaussian noise. Let \mathbf{s} be the vector of performance scores output by the original performance measurement. The new performance scores output by the manipulation-robust performance measurement are $\mathbf{s}' = \mathbf{s} + \boldsymbol{\epsilon}$, where every term of the vector $\boldsymbol{\epsilon}$ is drawn i.i.d. from a Gaussian distribution $g_\epsilon = \mathcal{N}(0, \sigma_\epsilon)$. Then, the mechanism rewards agents by applying a rank-order payment function on \mathbf{s}' . A manipulation-robust payment mechanism consists of a manipulation-robust performance measurement and a rank-order payment function.

Proposition 5.3. *For any payment mechanism consisting of a strongly truthful performance measurement and a winner-take-all tournament with $n \geq 2$ agents, there exists a threshold value, $\bar{\sigma}$, such that if the standard deviation of the noise is $\sigma_\epsilon > \bar{\sigma}$, the corresponding manipulation-robust payment mechanism is strongly truthful.*

We defer the proof to Appendix B.2. At a high level, the proof works because in terms of the expected payment of the untruthful deviation, adding a large noise weakens the advantage of the untruthful deviation (having a larger variance) while it enlarges the disadvantage (having a smaller expected score). the tradeoff between the gain from enlarging the variance and the detriment from decreasing the mean. We emphasize that the effectiveness of manipulation-robust payment mechanisms is premised on the assumption that the initial performance measurement is strongly truthful.¹³ Otherwise, untruthful deviations may increase the expected score or increase the variance without decreasing the mean of the performance score, in which cases adding noise cannot guarantee truthfulness. We further note two limitations of Proposition 5.3.

First, the current proof only applies to winner-take-all tournaments. The main reason for this is that the probability of being ranked first after deviation is monotonically decreasing with respect to σ_ϵ , which is not necessarily true for the probability of being ranked in any other position. However, we note that empirically, the idea works well for the other monotone rank-order payment functions considered in this paper.

Second, in theory, the required noise may have to be very large to guarantee the truthfulness of the manipulation-robust payment mechanism. As we will see in the next section, this is bad news. However, our empirical results suggest that the common noise need not to be large in practice. A reasonably small noise (e.g. the standard deviation of the noise is five times of the standard deviation of the original performance score) is enough guarantee truthfulness in many cases.

We further emphasize that adding noise to the performance score will not change any results in Section 4 as all the proofs trivially generalize.

5.3 The Variational Robustness

We demonstrate that although adding noise to the performance score helps guarantee the truthfulness of the manipulation-robust payment mechanism, it decreases the sensitivity of the original performance measurement. Intuitively, as the sensitivity, represented by $\delta = \frac{\mu'}{\sigma}$, is inversely proportional to the variance of the performance score, the added common noise will decrease the sensitivity, resulting in an increase in the total payment to incentivize a desired effort.

Proposition 5.4. *The sensitivity of the manipulation-robust performance measurement is decreasing in σ_ϵ , the standard deviation of the added noise.*

The proof straightforwardly follows as adding noise does not affect the numerator of the sensitivity while it increases the denominator. Proposition 5.4 suggests that under the rank-order payment function, there is a tradeoff between guaranteeing truthfulness and incentivizing a desired effort with a low cost of budget. Adding noise increases the robustness of the mechanism against strategic reporting but may require in a larger payment to elicit the goal effort. Therefore, while guaranteeing truthfulness, we want the variance of the added noise to be as small as possible.

Our discussions lead to a new aspect of the strategic robustness of the performance measurement. Under the tournament setting, a truthful performance measurement, which can punish any untruthful deviation by decreasing its expected score, is not enough. A robust performance measurement should also prevent

¹³However, the requirement of strongly truthfulness is not necessary. If a performance measurement is truthful, but it can guarantee no unilateral deviation can increase the variance without decreasing the mean of the performance score, the manipulation-robust payment mechanism can still guarantee truthfulness.

untruthful strategies from increasing the variance of the performance score. We name this property of a performance measurement the *variational robustness*, as we will define formally soon.

The concept of variational robustness is important as it relates to both the truthfulness of the payment mechanism and its ability to efficiently elicit a goal effort at a low cost. As explained in the previous section, ensuring the truthfulness of a payment mechanism may require adding noise on the performance score, which can decrease the sensitivity of the performance measurement. Therefore, performance measurements with lower variational robustness will have to scarify more of their sensitivity in order to achieve the truthfulness of the corresponding manipulation-robust payment mechanisms.

Definition 5.5. Given a strongly truthful performance measurement ψ and a fixed effort level ξ , let σ_τ be the standard deviation of the score distribution at the truth-telling strategy profile. Let σ_ϵ be the minimum standard deviation of the common noise that makes the manipulation-robust payment mechanism consisting of ψ and the winner-take-all RO-payment function truthful. Then, the variational robustness of ψ (at the effort level ξ) is defined as $\vartheta_\psi = \frac{\sigma_\tau}{\sqrt{\sigma_\tau^2 + \sigma_\epsilon^2}}$.

The variational robustness is defined as the ratio between the standard deviation of the truth-telling score distribution of the original performance measurement and that of the manipulation-robust performance measurement with the minimal noise, which takes a value from $(0, 1]$. A value of 1 indicates that under performance measurement ψ , any unilateral untruthful deviation cannot improve the expected payment of the payment mechanism consisting of ψ and the winner-take-all payment function, even without adding noise. It is worth noting that although ϑ is defined for truthful performance measurement under the winner-take-all RO-payment function, the concept can also be generalized to untruthful performance measurement and other monotone RO-payment functions in a straightforward manner. We will empirically evaluate this property of several commonly used performance measurements in Section 8.2.

6 Agent-based Model Setup and Assumption Justification

In this section, we describe the real-data estimated agent-based model we use for experiments. We examine the assumptions made for your theory. We show that either the assumption empirically holds, or the conclusion of our theory hold even if the assumptions do not.

6.1 Experiment Setup

6.1.1 Datasets

We use two crowdsourcing datasets to estimate the prior of ground truth w and agents' signal matrix Γ , called world 1 ($W1$) [3] and world 2 ($W2$) [30] respectively.

World 1 has a signal space with a size of five and a binary ground truth space, $\{1, 2\}$. Agents are asked to grade the synthetic accessibility of compounds with scores 1 to 5, where 1 indicates inappropriate to be synthesized and 5 stands for appropriate. Scores in between lower the confidence of the grading. The binary ground truth indicates whether a compound is appropriate or inappropriate. The dataset includes the assessments of 100 compound (tasks) from 18 agents. World 2 has a signal space and ground truth space which are both of size four.¹⁴ The dataset contains 6000 classifications of the sentiment of 300 tweets (tasks) provided by 110 workers. The estimated parameters for $W1$ and $W2$ are:

$$w_1 = [0.613 \ 0.387], \Gamma_1 = \begin{bmatrix} 0.684 & 0.221 & 0.032 & 0.037 & 0.026 \\ 0.092 & 0.191 & 0.050 & 0.200 & 0.467 \end{bmatrix};$$

$$w_2 = [0.196 \ 0.241 \ 0.247 \ 0.316], \Gamma_2 = \begin{bmatrix} 0.770 & 0.122 & 0.084 & 0.024 \\ 0.091 & 0.735 & 0.130 & 0.044 \\ 0.033 & 0.062 & 0.866 & 0.039 \\ 0.068 & 0.164 & 0.099 & 0.669 \end{bmatrix}.$$

Note that we use the estimated confusion matrices as the underlying full-effort working matrices Γ_{work} , which assumes the real-world agents are exerting full effort. Obviously, this is an under-estimation of Γ_{work} since the real-world agents' effort may be smaller than 1. Since the agents' effort cannot be directly observed, it is impossible to estimate Γ_{work} with no bias. However, it is not central to our experimental results depends on the correct estimation of Γ_{work} . Furthermore, the experiments are run with two different world models to show the robustness of our results.

¹⁴there are actually 5 signals, but we ignore the rarest one which only occurs 9 out of 300 times

6.1.2 The Performance Measurements

We implement two types of performance measurements: spot-checking and peer prediction mechanisms.

Spot-checking. Let p_c be the probability of spot-checking, i.e. the principal has the access to the ground truth of $n_c = p_c \cdot n$ randomly sampled tasks. We consider two spot-checking mechanisms in our experiments. First, a straightforward idea is to set the performance score to be the accuracy of the each agent’s reports on the spot-check questions. We denote this performance measurement as **SC-Acc**.

Alternatively, we apply the mechanism considered in [12], which is inspired by the Dasgupta-Ghosh mechanism [4]. We denote this performance measurement as **SC-DG**. Given an agent’s reports and a set of spot-checking questions with the ground truth, SC-DG randomly chooses a common task (bonus task) and two distinct tasks (penalty tasks). Then, the agent is scored 1 if her report on the bonus task agrees with the ground truth, and scored -1 if her report on the penalty task agrees with the ground truth of the distinct penalty task. The final score of the agent is the average score after repeated sampling.

Peer prediction. We consider five types of commonly used peer prediction mechanisms. The idea of peer prediction is to score each agent using some form of the correlation between her reports and her peers’ reports.

First, we implement the naive idea of paying an agent 1 if her report on a random task agrees with a random peer’s report on the same task, and paying 0 otherwise. This performance measurement is called the *output agreement mechanism* (**OA**) as discussed in [10].

Second, in the same paper, Faltings et al. [10] propose the *peer truth serum* (**PTS**) mechanism. The only difference between PTS and OA is that the payment is proportional to $\frac{1}{R(x)}$ when the pair of agents agree on a task, where R is a public distribution of reports and x is the report of those agents. While computing agent i ’s payment, we implement PTS by setting R to be the empirical distribution of all agents’ reports other than i .

Third, we consider the *matrix f -mutual information mechanism* (**f -MMI**). Inspired by Kong and Schoenebeck [22], **f -MMI** scores each agent using the estimation of the f -mutual information between her reports and her peer’s report, where f can be any convex function.

The f -MMI uses the empirical distributions to estimate the mutual information. We use the empirical distributions between two agents’ reports, i.e. $\tilde{P}_{\hat{X}_i, \hat{X}_j}$ for the joint distribution and $\tilde{P}_{\hat{X}_i}$ for the marginal distribution. Then, the MI between reports \hat{X}_i and \hat{X}_j can be estimated,

$$\widetilde{MI}_{i,j}^{f-MMI} = \sum_{x,y} \tilde{P}_{\hat{X}_i, \hat{X}_j}(x,y) f\left(\frac{\tilde{P}_{\hat{X}_i}(x)\tilde{P}_{\hat{X}_j}(y)}{\tilde{P}_{\hat{X}_i, \hat{X}_j}(x,y)}\right). \quad (5)$$

The matrix mutual information mechanism then scores each agent i using the average of the estimated MI between i and each of her peers. To speed up the mechanism, instead of pairing agent i with each of her peers, we simply learn the empirical distributions of the reports on each task of all agents but i . This can be seen as a “virtual agent” reporting based on the empirical distributions of all agents but i . Then, we learn the joint distribution as well as the mutual information between agent i ’s reports and this virtual agent’s reports.

Fourth, we implement the *pairing f -mutual information mechanism* (**f -PMI**) [26]. Similar to SC-DG, f -PMI randomly samples the bonus and penalty tasks and scores each agent based on whether her reports agree with the “ground truth” on the three tasks. The main difference is that instead of using the ground truth, the f -PMI learns a soft predictor on each task using all the other agents’ reports. Then, the f -mutual information is estimated for each agent using the soft-predictor and the agent’s reports. Note that the f -PMI contains the well known DG mechanism [4] and CA mechanism [28] as special cases when f is $f(x) = \frac{1}{2}|x - 1|$.

The f -PMI provides an alternative way to estimate the MI [26]. Specifically, the quotient of the joint distribution between \hat{X}_i and \hat{X}_{-i} and the product of the marginal can be written as

$$\frac{P_{\hat{X}_i, \hat{X}_{-i}}(\hat{x}_i, \hat{x}_{-i})}{P_{\hat{X}_i}(\hat{x}_i)P_{\hat{X}_{-i}}(\hat{x}_{-i})} = \frac{P_{\hat{X}_i|\hat{X}_{-i}}(\hat{x}_i|\hat{x}_{-i})}{P_{\hat{X}_i}(\hat{x}_i)}. \quad (6)$$

The denominator can be empirically estimated. While the numerator is estimated by the output of a soft-predictor, which, given the reports of all agents except i on a particular task j , produces a forecast of agent

i 's report on the same task in the form of a distribution. In our experiments, we set the soft-predictor for agent i 's report on task j as the empirical distribution of all the other agents' reports on the same task.

For both f -MMI and f -PMI, we consider four types of commonly used f -divergence for the MMI and PMI mechanisms, as shown in Table 1.

Finally, we implement the *determinant mutual information mechanism (DMI)* [19]. Kong generalizes the Shannon mutual information to the determinant mutual information. Specifically, for a pair of agents i and j , the set of the commonly answered tasks is divided into two disjoint subsets A and B . Again, we empirically estimate the joint distribution with reports in A and B respectively, and score agent i using the product of the determinants of these two estimated joint distribution matrices. Finally, the score of each agent is determined by averaging the scores obtained from pairing that agent with every other agent.

Table 1: Four f -divergences

f -divergence (short name)	$f(a)$
Total variation distance (TVD)	$\frac{1}{2} a - 1 $
KL-divergence (KL)	$a \log a$
Pearson χ^2 (Sqr)	$(a - 1)^2$
Squared Hellinger (Hlg)	$(1 - \sqrt{a})^2$

6.1.3 Parameters And Estimation Methods

Unless otherwise specified, we set the number of tasks to be $m = 1000$ with each agent answering $m_a = 100$ tasks. Every task is assigned to (at least) $n_0 = 5$ agents and there are $n = 52$ agents in total.¹⁵ We consider two types of commonly used cost functions: polynomial cost $c(e) = e^r$ and exponential cost $c(e) = \exp(r \cdot e)$. When dealing with loss-averse agents, ρ is set to be 0.5. For spot-checking mechanisms, we vary the spot-checking probability from 0.1 to 0.3. For peer prediction mechanisms, we use four types of commonly used f -mutual information with f listed in Table 1.

While considering deviations in effort levels, for each performance measurement, we estimate the distributions of the performance score of an agent before and after a unilateral deviation of effort $\xi + \Delta e$ for $\xi \in \{0, 0.01, \dots, 0.99\}$. Fixing each of the ξ , we first simulate the report matrix \mathbf{x} when all agents exert an effort of ξ . Then, we input \mathbf{x} to each performance measurement, which gives us n samples (one for each agent) of the performance score before deviation. Then, let one of the n agents deviates to an effort level of $\xi + \Delta e$ with $\Delta e = 0.01$. Repeating the above process gives us one sample of the performance score after the unilateral deviation. We further repeat the process so as to generate 5000 samples for the performance score generated by each performance measurement, at each effort level, before and after deviation. Finally, we fit the Gaussian models with the generated samples. By estimating the mean and the standard deviation, we then can estimate the probability of each rank. In this way, the optimal RO-payment function can be developed based on our theoretical results from Section 4.1.¹⁶

While considering agents' reporting strategies, we generate samples and estimate the Gaussian model in the same way as above. In particular, fixing an effort ξ , we generate the samples of the performance score of three cases: 1) all agents are truthful; 2) an agent i 's performance score when she deviates to an untruthful strategy π ; 3) and other agents' performance scores when agent i deviates.¹⁷ Then, we fit the samples to the Gaussian models. We name the estimated distribution \hat{g}_i , with the corresponding mean and standard deviation μ_i and θ_i for $i = 1, 2, 3$, respectively.

For the space of untruthful strategies, we heuristically choose a large set of strategies that merge one signal to another. Specifically, while seeing a signal s , the agents report $\pi(s) \neq s$ with some probability (fixed at 0.5 in our experiments).¹⁸ For example, three types of the strategies in $W1$ can be (1) mapping signal $4 \rightarrow 5$, (2)

¹⁵The number of agents $n > 50$ is to guarantee that each task is assigned with at least $n_0 = 5$ agents and tasks are assigned to agents randomly.

¹⁶Note that when n is finite, there may not be integer solutions for the thresholds \bar{n} and \hat{n} in our propositions, in which cases the thresholds are rounded to the closest integers.

¹⁷As we will see in Section 6.2, the score distribution of case 3) is almost identical to that of case 1). In the following sections, we thus use the estimated score distribution of case 1) for the truth-telling distribution.

¹⁸The main reason that we consider mixed strategy is to avoid missing signals in agents' reports, which may greatly decrease the variance of performance scores. For example, it trivially results in scores of zeros, and thus results in a variance of zero for

mapping $1 \rightarrow 2$ and $5 \rightarrow 4$ and (3) mapping $x \rightarrow x - 1$ for $x \in \{2, 3, 4, 5\}$.

6.2 Assumption Justifications

In our theoretical analysis, we made several assumptions to support our conclusions. To ensure the validity of our theoretical results, we now conduct empirical verification through agent-based model (ABM) experiments and demonstrate that the assumptions made in our theory (approximately) capture the characteristics of the real problem.

First, we assume the distribution of agents' performance score follows a Gaussian distribution. We note that this assumption is obviously violated if the performance measurements have bounded scores, such as the matrix mutual information mechanism (MMI) and both of the spot-checking mechanisms. However, our experiments show that the Gaussian distribution can approximate the performance score distributions for most of the considered performance measurements. In Fig. 3, we present the Kolmogorov-Smirnov test on the empirical distribution (estimated with 5200 samples) and the estimated Gaussian distribution for each of the considered performance measurement [29]. A higher KS test statistic indicates a larger discrepancy between the two distributions and hence a higher error in the Gaussian estimation. Note that a KS test statistic of 0.02 is equivalent to the KS test statistic between a Gaussian distribution and the empirical distribution estimated by 2000 i.i.d. samples from that Gaussian distribution; while the number of i.i.d. samples to achieve a KS test statistic of 0.05 is about 500. From Fig. 3, we observe that most of the performance measurements can be fitted reasonably well, except DMI, *Sqr*-MMI, *KL*-PMI and *Hlg*-PMI, whose score distributions tend to have heavy-tailed (as shown in Fig. 4(b)) and thus are not well-fitted by the Gaussian model. However, exactly because of these heavy tails, these performance measurements have high variances and thus low sensitivities, which are unlikely to perform well in realistic settings. Furthermore, to motivate Assumption 4.1, we observe that while changing the effort, the change of the standard deviation of the performance score is usually less than 0.1 times the change of the mean. Take Fig. 4(a) as an example, where the mean of the performance score changes by 0.054 after deviation, while the std changes by -0.005 .


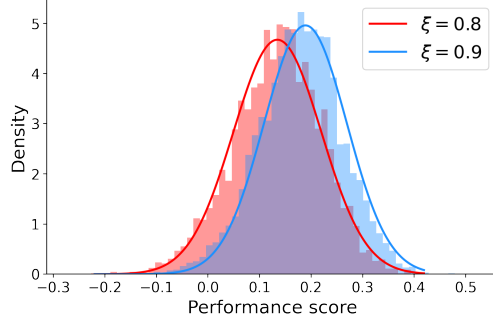


Figure 3: The Kolmogorov-Smirnov test statistics for different performance measurements in $W1$. Each data point is computed by applying the KS test between the empirical distribution estimated with 5200 samples and the Gaussian distribution estimated using the same set of samples.


Second, in Assumption 3.2 and 3.5, we assume that unilateral deviations both in effort and in reporting strategies will not significantly affect the performance score distributions of other agents. However, these assumptions may not hold under peer prediction mechanisms where an agent's performance score may depend on other agents' strategies, especially when the number of agents is not large enough. To verify the validity of these assumptions, we conduct experiments to estimate and compare the score distributions of other agents before and after one agent's unilateral deviation. We observe that the change in the estimated distributions is minimal even for relatively small groups of agents in the crowdsourcing setting, e.g. $n = 50$.

Third, as is common, we assume that the first order condition (FOC) is sufficient for the existence and uniqueness of equilibrium. With our ABM experiments, we find that this assumption is valid for the considered performance measurements, RO-payment functions and cost functions. Specifically, we observe that the

DMI.



(a) For *TVD-PMI*, Gaussian is a good fit.



(b) For *KL-PMI*, Gaussian is not a good fit.

Figure 4: The histogram of the empirical distributions of the performance score v.s. the fitted Gaussian distributions. Both examples are in $W1$.

expected utility (as defined in Eq. (3)) is concave w.r.t. e_i in our settings, which implies that there exists a unique e_i that maximizes each agent’s expected utility. Therefore, FOC is sufficient for both the existence and uniqueness of the symmetric equilibrium.

Finally, while Lemma 4.2 is derived under the assumption of an infinite number of agents, our experiments demonstrate that the lemma holds even for a reasonable number of agents in the crowdsourcing setting, e.g. $n = 50$. This means that our analytical results of the optimal RO-payment functions shown in Section 4 still apply even when the number of agents is not extremely large.

7 An Agent-based Analysis of RO-Payment Functions

Although the winner-take-all tournament is shown to be optimal in our setting (Proposition 4.3) and several similar principal-agent settings [23, 8], it is not commonly used in practice. In Section 4, we show that this may be due to agents’ loss-aversion and risk-aversion, or the need to compensate for their costs of effort to ensure individual rationality (IR). In this section, we use our agent-based model to take an in-depth analysis of how the inclusiveness of the optimal RO-payment function relates to the model parameters, including the goal effort level, agents’ cost functions and utility models.

Furthermore, we empirically compare the rank-order payment function with the linear payment function. We show that even when accounting for the noise added to guarantee truthfulness, the optimized rank-order payment function is still significantly more budget efficient than the linear payment function, which guarantees truthfulness without the need of adding noise.

7.1 Inclusiveness and Model Parameters

Using our agent-based model simulations, we visualize the inclusiveness of the optimal RO-payment function, which indicates the number of agents that receive non-zero payments. Figure 5 presents examples of the inclusiveness of the optimal RO-payment functions under three utility models, while varying the goal effort ξ and the cost function c .

Our first observation is that the IR constraint is likely to be binding when the cost function is “less convex” and the effort ξ is high. This observation is in line with our theory as whether IR is binding depends on the ratio $\frac{c'(\xi)}{c(\xi)}$. Taking Theorem 4.3 as an example, IR is not binding when $\frac{c'(\xi)}{c(\xi)} > np'_1(\xi)$. For most commonly used cost functions, such as quadratic or higher order polynomial functions and exponential functions, this ratio is likely to be larger with more convex functions and with larger ξ . In both cases, compensating agents’ costs becomes more expensive, resulting in IR-minimal payments and higher inclusiveness of the optimal RO-payment functions.

Furthermore, by comparing lines of the same color across different utility models, we observe that in general, when agents are loss-averse or risk-averse, the inclusiveness of the optimal RO-payment function is higher than when agents are neutral, especially when the goal effort is high. This finding is consistent with our conclusion from Section 4.1.2.

In practice, it is common for crowdsourcing agents to be risk-averse and loss-averse, and because tasks such as labeling are usually not extremely difficult, a higher effort does not result in a significantly larger marginal cost, meaning that the cost functions are not highly convex. Therefore, our findings in this section suggest that inclusive payment functions are likely to be preferred in realistic crowdsourcing settings, given these characteristics of agents and tasks.




Figure 5: The inclusiveness of the optimal RO-payment functions under different agent utilities as a function of the symmetric equilibrium effort ξ with different cost functions. The solid curves represent the inclusiveness of the optimal RO-payment functions when considering IR, while the dashed curves show the inclusiveness when IR is ignored. Note that for (a) and (b), the optimal RO-payment function does not depend on the cost functions and is represented by a single grey dashed curve. In this figure, we use SC-Acc as the performance measurement with the spot-checking probability of 0.25. For (b), $\rho = 0.5$ and for (c), $r_a(t) = \log(t + 1)$.


7.2 Rank-Order versus Linear Payment Functions

Recall that although linear payment functions can preserve the truthfulness of the payment mechanism, they are not flexible to optimize and may experience an enormous budgetary cost in many cases. Here, we provide a more fine-grained comparison between the payment of the linear function and the payment of the rank-order function, while in both cases, we require the payment mechanism satisfying limited liability (1), effort elicitation (3) and truthfulness (4).

Parameters of Linear Payment Functions. A linear payment function rewards an agent $t_i = a \cdot s_i + b$ where s_i is the performance score and a and b are constants. To incentivize a certain effort level, ξ , as an equilibrium, we set the factor a to $\frac{c'(\xi)}{\mu'(\xi)}$, where $c'(\xi)$ is the derivative of the cost function at the goal effort ξ , and $\mu'(\xi)$ is the derivative of the expected performance score w.r.t. an agent’s effort when every agent’s effort is ξ . We then set b to satisfy limited liability. However, for performance measurements with unbounded performance scores, no constant factor b can guarantee limited liability. To address this issue, we modify the linear payment function by treating b as a variable, denoted as \tilde{b} , that is computed such that the minimum payment is equal to zero. This modification makes \tilde{b} depend on agents’ reporting strategies and thus does not technically preserve the truthfulness of the performance measurement.¹⁹ However, \tilde{b} is a lower bound of b . Therefore, the payment of modified linear payment function lower bounds the payment of the “real” payment function. We will show that even compared with its lower bound, the RO-payment function induces much smaller payments than the linear payment function.

Adding Noise to RO-payment Functions. Next, we apply the idea of the manipulation-robust payment mechanism to make rank-order payments truthful. For every goal effort and every untruthful deviation, we empirically compute the minimum variance of the noise that can guarantee truthfulness. Then, for every goal effort, we pick the largest required noise which guarantees that no unilateral deviation (in the strategy space that we consider) will result in a larger expected payment. The added noise enlarges the variance of the score distributions which decreases the sensitivity, and thus increases the minimum payment to incentivize the goal effort. Our comparison is between the modified linear payment function discussed above and the manipulation-robust RO-payment function after adding the noise.

¹⁹While it is theoretically possible to game such a modified linear payment function with a unilateral untruthful deviation, it is highly unlikely to be successful. The potential benefit from such a deviation is limited, as the agent would have to decrease the



(a) Matrix mutual information mechanism (b) Pairing mutual information mechanism (c) Spot-checking with accuracy score (SC- with the Hellinger divergence (Hlg -MMI). with the Hellinger divergence (Hlg -PMI). Acc).

Figure 6: The comparison between the total payments of the modified linear payment functions and the manipulation-robust rank-order payment functions. All three examples are in the case of risk-neutral agents (and thus the corresponding optimal RO-payment function when IR is not binding is winner-take-all by Theorem 4.3), and use the cost function of $c(e) = e^2$. The dashed curves in three examples correspond to the IR-minimal payment which equals to $n \cdot c(\xi)$, and thus are identical in all three figures.

Results. In Fig. 6, we compare the total payment of the modified linear payment function with the manipulation-robust RO-payment function. Our first observation is that the payments of the modified linear payment function (which are lower bounds of the actual required payments to preserve truthfulness under the linear payment function) are much larger than the payments of the RO-payment functions. This observation consistently holds for every goal effort, and is particularly pronounced for the performance measurements whose performance scores are unbounded below (e.g. the KL -PMI shown in Fig. 6 (b)), where the payments of the linear function are hundreds or even thousands times of the IR-minimal payment. However, even for performance measurements which have non-negative performance scores (e.g. the Hlg -MMI shown in Fig. 6 (a)) where factor b can be zero or even negative, the linear payment function is still significantly dominated by the optimal RO-payment function.

The second takeaway is that the RO-payment function is very effective in eliciting the goal effort, which can achieve the IR-minimal payment for a large range of goal effort. This can be observed from the figures where the solid curves are almost identical to the black dashed curves.²⁰

We note that although the examples in Fig. 6 are based on the winner-take-all tournament, similar pattern can be observed while considering more inclusive RO-payment functions. Furthermore, we observe that a more inclusive RO-payment function is more robust against strategic reporting. That is, with a more inclusive payment function, a deviating agent needs a larger increase in the variance of the performance score to gain an advantage, which in turn reduces the amount of noise needed.

Our observations warn that linear payment functions may not be practical in real-world scenarios. When budget efficiency is a big concern, we note that the rank-order payment function is a good choice.

8 Evaluating Realistic Performance Measurements

In the idealized setting, we have reduced the optimization of performance measurements to the problem of maximizing the sensitivity of a performance measurement. However, in practice, we cannot arbitrarily increase the sensitivity of a performance measurement as desired, but are given a limited set of options such as spot-checking and peer prediction mechanisms. We then ask: which mechanism has the highest sensitivity and how does the sensitivity changes as the goal effort level increases.

Additionally, achieving a high sensitivity is only one part of the problem. To ensure truthfulness under the rank-order payment function, noise must be added to discourage strategies that would benefit from increasing the variance of the performance score. As shown in Section 5.3, this noise can negatively impact sensitivity. Therefore, the variational robustness of a performance measurement serves as a second dimension of our problem.

In this section, we use our agent-based model to empirically evaluate several commonly considered

score of the bottom agent more than the decrease in her own score, which is unlikely to occur when the number of agents is large.

²⁰With an exception of the pairing mechanism, which is less robust against strategic reporting and thus requires a larger noise to maintain truthfulness.

performance measurements as introduced in Section 6.1.2 with respect to their sensitivity and variational robustness. Based on this comparison, we will recommend the best mechanism(s) for use.

8.1 The Sensitivity of Performance Measurements

In Fig. 7, we show the (smoothed) sensitivity versus the goal effort for different performance measurements, where higher curves are preferred. Note that there are cases where the estimated sensitivity can be negative. This is because when the sensitivity is close to zero, the estimate may veer negative due to the limited number of samples. Our results are summarized below:

1. The spot-checking mechanisms generally have consistent sensitivities, while most of the peer prediction mechanisms have increasing sensitivities. This is because as the goal effort increases, the reports of agents become more accurate, allowing peer prediction mechanisms to estimate the correlations between agents' reports more accurately and thus improving their sensitivities. However, since the performance score of the spot-checking mechanism does not depend on the peers' reports and effort, their sensitivities remain consistent across different effort levels.
2. It is not surprising that the sensitivity of spot-checking mechanisms increases as the spot-checking probability increase, as more spot-checks lead to a more accurate estimation of agents' effort. Furthermore, there is no significant difference between SC-Acc and SC-DG.
3. The sensitivities of peer prediction mechanisms vary significantly. In general, if we focus on the high-effort range with $\xi \geq 0.6$, OA, a simple scoring rule has the highest sensitivity. The f -MMI mechanisms also perform well on both datasets, while the f -PMI mechanisms are less sensitive.

In summary, when the goal effort is low, our experiments suggest the use of spot-checking. When the goal effort is large, peer prediction based performance measurements are better choices as they are more sensitive and cheaper to apply. In terms of budget efficiency, the best performance measurements are SC-Acc for the spot-checking based performance measurement, and OA, Hlg -MMI or KL -MMI for peer prediction based performance measurement.

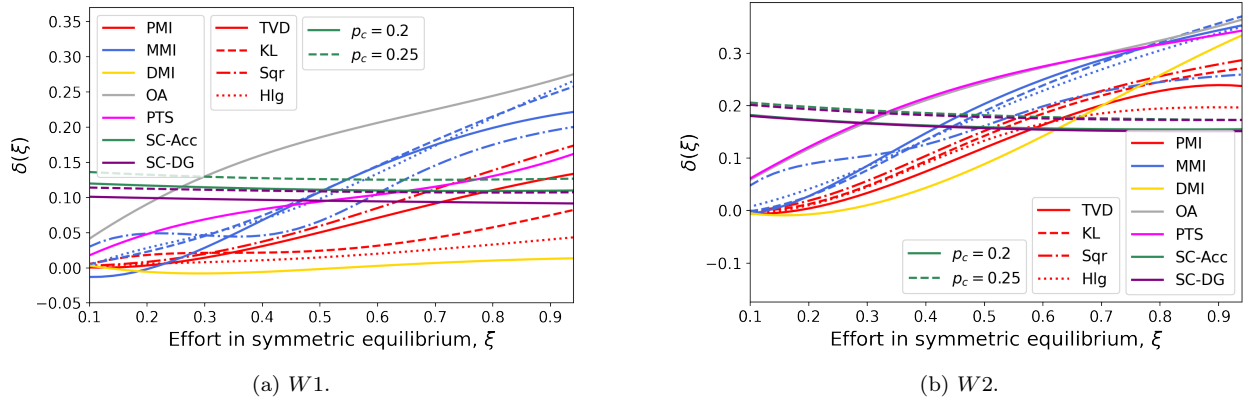


Figure 7: The sensitivity $\delta(\xi)$ versus the goal effort ξ for different performance measurements in (a) W1 and (b) W2.

8.2 The Variational Robustness of Performance Measurements

Recall that the variational robustness of a performance measurement (see Definition 5.5), is defined based on two standard deviations: the std of the score distribution under the truth-telling strategy profile and the std of the minimum size of noise required to guarantee truthfulness. Using the same method as in Section 6.1.3, we can estimate both of the standard deviations, which provides an estimation of the variational robustness.

It is worth noting that our comparisons include some performance measurements that are not (strongly) truthful, namely, OA, PTS and spot-checking mechanisms.²¹ For these performance measurements, it is

²¹Spot-checking mechanisms are not truthful when the ground truth space is smaller than the signal space (for example, in W1).

possible that some untruthful strategies can increase the expected score, in which case adding noise will not preserve the truthfulness of the original performance measurements. In order to present a comparison of the variational robustness of different performance measurements, we ignore the strategies that increase the expected score while estimating their variational robustness. However, we emphasize that for the mechanisms that are not truthful, even a variational robustness of 1 does not imply that they can penalize all untruthful unilateral deviations.

In Fig. 8, we present the results of our comparisons of sensitivity and variational robustness for various performance measurements. Although the variational robustness depends on the goal effort ξ and the RO-payment functions, the following patterns generally hold.

First, the output agreement mechanism (OA), though it is not truthful, tends to have high variational robustness. This means untruthful deviations may increase the expected score, but it is hard for them to greatly increase the variance of the performance score. This property can be attributed to the simplicity of OA. Similarly, with the help of the ground truth information, the spot-checking mechanisms also have high variational robustness.

Second, for truthful performance measurements, the pairing mutual information mechanisms (PMI) are not variationally robust, while the matrix mutual information mechanisms (MMI), particularly *KL*-MMI and *Hlg*-MMI, are consistently robust when the goal effort is high, e.g. $\xi \geq 0.5$.

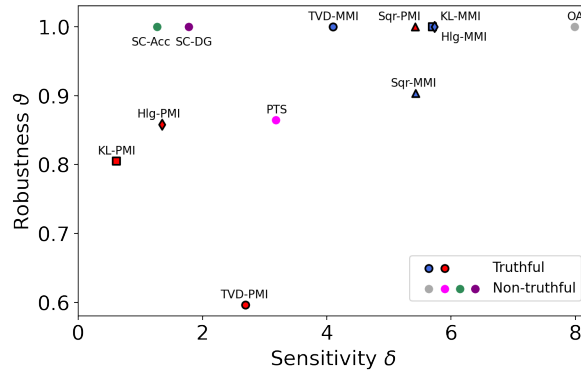


Figure 8: The sensitivity and the variational robustness of difference performance measurements in *W1*. The goal effort is fixed at $\xi = 0.8$. Performance measurements that are theoretically truthful have markers with black edges while those are not truthful have no edge.²²

In summary, our results suggest that if truthfulness is not a primary concern, the output agreement mechanism (OA) is a good choice due to its simplicity and high sensitivity. While if we additionally want the payment mechanism to be truthful, the matrix mutual information mechanisms *KL*-MMI and *Hlg*-MMI are strong candidates as they have high sensitivity and high variational robustness and are (approximately) truthful. Additionally, if some ground truth information is available, the spot-checking mechanism can be a good option as they are truthful (if the ground truth space equals the signal space), have high sensitivity and high variational robustness, especially when the goal effort is low.

9 Conclusion and Future Work

We study the problem of how to simultaneously incentivize effort and elicit truthful reporting when agents are strategic in their choice of effort level and reporting strategy. To do so, we propose a two-stage framework of payment mechanism design, compositing of a performance measurement and a rank-order payment function, which addresses practical concerns including limited liability and budget efficiency. Our main contributions are:

- We establish four objectives of mechanism design in crowdsourcing, and show that previous approaches can only simultaneously achieve three at most.

²²Note that the matrix mutual information mechanisms (MMI) are actually approximately truthful (a slightly weaker version of truthfulness) and the error vanishes as m , the number of tasks, is large enough.

- Our payment mechanism framework is the first that puts forth the study of peer prediction mechanisms into the setting of continuous effort.
- We optimize the rank-order payment function to incentivize a goal effort while minimizing the cost of budget. We fill a gap in the principal-agent problem. by providing analytical solutions to the optimal rank-order payment functions when individual rationality is a hard constraint.
- We identify a sufficient statistic for evaluating the budget efficiency of a performance measurement.
- We propose an idea of adding noise to agents' performance scores to preserve the truthfulness of a performance measurement under the non-linear rank-order payment function, which helps us simultaneously achieve all four objectives.
- Our agent-based model experiments evaluate commonly used performance measurements and provide practical guidance for their implementations.

Several promising future directions exist. First, heterogeneous agents that have different cost functions and confusion matrices could serve as a potential generalization of this paper, where asymmetric equilibrium may be considered. Second, we focus on rank-based payments in this paper, but some of our insights can be generalized to other contracts, e.g. the independent contract[14]. The problems of which contract is optimal under which circumstances still remains open. Finally, although the rank-order payment functions do not require much information from the principal, they do require some. In particular, at the desired effort, the agents' cost and its derivative, and agents' signal distributions must be estimated. How can these parameters be learned by the principal and how robust are mechanisms to misspecifications of these parameters?

References

- [1] Arpit Agarwal, Debmalya Mandal, David C Parkes, and Nisarg Shah. 2017. Peer Prediction with Heterogeneous Users. In *Proceedings of the 2017 ACM Conference on Economics and Computation*. ACM, 81–98.
- [2] Mara S. Aruguete, Ho Huynh, Blaine L. Browne, Bethany Jurs, Emilia Flint, and Lynn E. McCutcheon. 2019. How serious is the ‘carelessness’ problem on Mechanical Turk? *International Journal of Social Research Methodology* 22, 5 (2019), 441–449. <https://doi.org/10.1080/13645579.2018.1563966> arXiv:<https://doi.org/10.1080/13645579.2018.1563966>
- [3] Isomura Tetsu. Baba Yukino, Kashima Hisashi. 2018. Data for: Wisdom of Crowds for Synthetic Accessibility Evaluation. *Mendeley Data* 1 (2018). <https://doi.org/10.17632/nmdz4yfk2t.1>
- [4] 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.
- [5] A. P. Dawid and A. M. Skene. 1979. Maximum Likelihood Estimation of Observer Error-Rates Using the EM Algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)* 28, 1 (1979), 20–28. <http://www.jstor.org/stable/2346806>
- [6] Mikhail Drugov and Dmitry Ryvkin. 2019. Optimal prizes in tournaments with risk-averse agents.
- [7] Mikhail Drugov and Dmitry Ryvkin. 2020. How noise affects effort in tournaments. *Journal of Economic Theory* 188 (2020), 105065. <https://doi.org/10.1016/j.jet.2020.105065>
- [8] Mikhail Drugov and Dmitry Ryvkin. 2020. Tournament rewards and heavy tails. *Journal of Economic Theory* 190 (2020), 105116. <https://doi.org/10.1016/j.jet.2020.105116>
- [9] David Easley and Arpita Ghosh. 2015. Behavioral Mechanism Design: Optimal Crowdsourcing Contracts and Prospect Theory. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation* (Portland, Oregon, USA) (*EC '15*). Association for Computing Machinery, New York, NY, USA, 679–696. <https://doi.org/10.1145/2764468.2764513>

- [10] Boi Faltings, Radu Jurca, and Goran Radanovic. 2017. Peer Truth Serum: Incentives for Crowdsourcing Measurements and Opinions. *arXiv:1704.05269* [cs.GT]
- [11] Rafael Frongillo, Robert Gomez, Anish Thilagar, and Bo Waggoner. 2021. Efficient Competitions and Online Learning with Strategic Forecasters. <https://doi.org/10.48550/ARXIV.2102.08358>
- [12] Alice Gao, James R. Wright, and Kevin Leyton-Brown. 2016. Incentivizing Evaluation via Limited Access to Ground Truth: Peer-Prediction Makes Things Worse. *arXiv:1606.07042* [cs.GT]
- [13] Arpita Ghosh and Patrick Hummel. 2012. Implementing Optimal Outcomes in Social Computing: A Game-Theoretic Approach. *arXiv:1202.3480* [cs.GT]
- [14] Jerry R. Green and Nancy L. Stokey. 1983. A Comparison of Tournaments and Contracts. *Journal of Political Economy* 91, 3 (1983), 349–364. <http://www.jstor.org/stable/1837093>
- [15] Chien-Ju Ho, Aleksandrs Slivkins, Siddharth Suri, and Jennifer Wortman Vaughan. 2015. Incentivizing High Quality Crowdsourcing. In *Proceedings of the 24th International Conference on World Wide Web* (Florence, Italy) (*WWW '15*). International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 419–429. <https://doi.org/10.1145/2736277.2741102>
- [16] Chien-Ju Ho, Aleksandrs Slivkins, and Jennifer Wortman Vaughan. 2015. Adaptive Contract Design for Crowdsourcing Markets: Bandit Algorithms for Repeated Principal-Agent Problems. *arXiv:1405.2875* [cs.DS]
- [17] Ajay Kalra and Mengze Shi. 2001. Designing Optimal Sales Contests: A Theoretical Perspective. *Marketing Science* 190, 2 (2001), 170–193. <https://doi.org/10.1287/mksc.20.2.170.10193>
- [18] Nur Kaynar and Auyon Siddiq. 2022. Estimating Effects of Incentive Contracts in Online Labor Platforms. *Management Science* (2022).
- [19] Yuqing Kong. 2020. Dominantly Truthful Multi-Task Peer Prediction with a Constant Number of Tasks. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms* (Salt Lake City, Utah) (*SODA '20*). Society for Industrial and Applied Mathematics, USA, 2398–2411.
- [20] Yuqing Kong and Grant Schoenebeck. 2018. Eliciting Expertise without Verification. In *Proceedings of the 2018 ACM Conference on Economics and Computation* (Ithaca, NY, USA) (*EC '18*). Association for Computing Machinery, New York, NY, USA, 195–212. <https://doi.org/10.1145/3219166.3219172>
- [21] Yuqing Kong and Grant Schoenebeck. 2019. An information theoretic framework for designing information elicitation mechanisms that reward truth-telling. *ACM Transactions on Economics and Computation (TEAC)* 7, 1 (2019), 2.
- [22] Yuqing Kong and Grant Schoenebeck. 2019. An Information Theoretic Framework For Designing Information Elicitation Mechanisms That Reward Truth-Telling. *ACM Trans. Econ. Comput.* 7, 1, Article 2 (Jan. 2019), 33 pages. <https://doi.org/10.1145/3296670>
- [23] Vijay B. Krishna and John Morgan. 1998. The Winner-Take-All Principle in Small Tournaments.
- [24] Barry J. Nalebuff and Joseph E. Stiglitz. 1983. Prizes and Incentives: Towards a General Theory of Compensation and Competition. *The Bell Journal of Economics* 14, 1 (1983), 21–43. <http://www.jstor.org/stable/3003535>
- [25] Truls Pedersen and Sjur Kristoffer Dyrkolbotn. 2013. Agents Homogeneous: A Procedurally Anonymous Semantics Characterizing the Homogeneous Fragment of ATL. In *PRIMA 2013: Principles and Practice of Multi-Agent Systems*, Guido Boella, Edith Elkind, Bastin Tony Roy Savarimuthu, Frank Dignum, and Martin K. Purvis (Eds.). Springer Berlin Heidelberg, Berlin, Heidelberg, 245–259.

- [26] Grant Schoenebeck and Fang-Yi Yu. 2021. Learning and Strongly Truthful Multi-Task Peer Prediction: A Variational Approach. In *12th Innovations in Theoretical Computer Science Conference (ITCS 2021) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 185)*, James R. Lee (Ed.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 78:1–78:20. <https://doi.org/10.4230/LIPIcs.ITCS.2021.78>
- [27] Grant Schoenebeck, Fang-Yi Yu, and Yichi Zhang. 2021. Information Elicitation from Rowdy Crowds. In *Proceedings of the Web Conference 2021 (Ljubljana, Slovenia) (WWW '21)*. Association for Computing Machinery, New York, NY, USA, 3974–3986. <https://doi.org/10.1145/3442381.3449840>
- [28] 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 (Maastricht, The Netherlands) (EC '16)*. ACM, New York, NY, USA, 179–196. <https://doi.org/10.1145/2940716.2940790>
- [29] Nickolay Smirnov. 1948. Table for estimating the goodness of fit of empirical distributions. *The annals of mathematical statistics* 19, 2 (1948), 279–281.
- [30] Matteo Venanzi, William Teacy, Alexander Rogers, and Nicholas Jennings. 2015. Weather Sentiment - Amazon Mechanical Turk dataset. <https://eprints.soton.ac.uk/376543/>
- [31] Wanyuan Wang, Bo An, and Yichuan Jiang. 2020. Optimal Spot-Checking for Improving the Evaluation Quality of Crowdsourcing: Application to Peer Grading Systems. *IEEE Transactions on Computational Social Systems* PP (06 2020), 1–16. <https://doi.org/10.1109/TCSS.2020.2998732>
- [32] Jens Witkowski, Rupert Freeman, Jennifer Wortman Vaughan, David M Pennock, and Andreas Krause. 2022. Incentive-compatible forecasting competitions. *Management Science* (2022).
- [33] Yichi Zhang and Grant Schoenebeck. 2023. Multitask Peer Prediction With Task-dependent Strategies. In *Proceedings of the Web Conference 2023 (Austin, TX, USA) (WWW '23)*. Association for Computing Machinery, New York, NY, USA. <https://doi.org/10.1145/3543507.3583292>

A Proofs and Details of Section 4

A.1 The Rank-order Impact After Convergence

Proof of Lemma 4.2. Fixing ξ , we simply let $g_e \sim N(\mu(e, \xi), \sigma(e, \xi))$ be the p.d.f. of the scores when agent i 's effort is e and all the other agents' effort is ξ , and let G_e be the c.d.f.. Let S be a random variable with p.d.f. g_e . Let $q_e(p)$ be the quantile function of S such that $\int_{-\infty}^{q(p)} g_e(x) dx = p$.

Because $p_j(\xi, \xi) = \frac{1}{n}$, it's equivalent to show that $p_j(\xi', \xi)$ is decreasing in j , where $\xi' = \xi + \Delta e$. Note that $p_j(\xi', \xi)$ is the j th order statics, which concentrates to its expectation when n is sufficiently large. Therefore, $p_j(\xi', \xi)$ can be approximated by the quantile function, i.e. $p_j(\xi', \xi) = G_{\xi'}(q_\xi(1 - \frac{j}{n})) - G_{\xi'}(q_\xi(1 - \frac{j+1}{n}))$. Let $\mu = \mu(\xi, \xi)$ and $\Delta\mu = \mu(\xi', \xi) - \mu(\xi, \xi)$. Let σ and $\Delta\sigma$ be the similar notations for std. Note that $\Delta e \rightarrow 0$ implies $\Delta\mu \rightarrow 0$ and $\Delta\sigma \rightarrow 0$ since $\mu(e)$ and $\sigma(e)$ are differentiable (Assumption 3.1).

We first prove the following intermediate step.

Lemma A.1. $G_{\xi'}(x) \approx (1 - \Delta\sigma/\sigma) G_\xi(x) - (\Delta\mu + \Delta\sigma)g_\xi(x)$.

Proof.

$$G_{\xi'}(x) = \frac{1}{\sqrt{2\pi}(\sigma + \Delta\sigma)} \int_{-\infty}^x e^{-\frac{1}{2}(\frac{s-\mu-\Delta\mu}{\sigma+\Delta\sigma})^2} ds$$

When $\Delta\sigma \ll \sigma$, $\frac{1}{\sigma+\Delta\sigma} = \frac{\sigma-\Delta\sigma}{\sigma^2-\Delta\sigma^2} \approx \frac{\sigma-\Delta\sigma}{\sigma^2} = \frac{1}{\sigma}(1 - \frac{\Delta\sigma}{\sigma})$. Therefore, $\frac{s-\mu-\Delta\mu}{\sigma+\Delta\sigma} \approx (1 - \frac{\Delta\sigma}{\sigma}) \frac{s-\mu}{\sigma} - \frac{\Delta\mu}{\sigma}$. We can rewrite the integrand by omitting the second-order infinitesimals.

$$\begin{aligned} &\approx \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^x \left(1 - \frac{\Delta\sigma}{\sigma}\right) e^{-\frac{1}{2}((1-\frac{\Delta\sigma}{\sigma})\frac{s-\mu}{\sigma} - \frac{\Delta\mu}{\sigma})^2} ds \\ &\approx \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^x \left(1 - \frac{\Delta\sigma}{\sigma}\right) e^{-\frac{1}{2}((1-\frac{\Delta\sigma}{\sigma})\frac{s-\mu}{\sigma})^2 + (1-\frac{\Delta\sigma}{\sigma})\frac{(s-\mu)}{\sigma}\frac{\Delta\mu}{\sigma}} ds \end{aligned}$$

By utilizing the Taylor expansion of e^x and disregarding higher-order infinitesimals, we can arrive at the approximation that $e^x \approx 1 + x$ when $x \rightarrow 0$. We apply this property on $e^{(1-\frac{\Delta\sigma}{\sigma})\frac{(s-\mu)}{\sigma}\frac{\Delta\mu}{\sigma}}$.

$$\begin{aligned} &\approx \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^x \left(1 - \frac{\Delta\sigma}{\sigma}\right) \left(1 + \left(1 - \frac{\Delta\sigma}{\sigma}\right) \frac{(s-\mu)}{\sigma} \frac{\Delta\mu}{\sigma}\right) e^{-\frac{1}{2}((1-\frac{\Delta\sigma}{\sigma})\frac{s-\mu}{\sigma})^2} ds \\ &\approx \frac{1}{\sqrt{2\pi}\sigma} \int_{-\infty}^x \left(1 - \frac{\Delta\sigma}{\sigma}\right) \left(1 + \frac{(s-\mu)}{\sigma} \frac{\Delta\mu}{\sigma}\right) e^{-\frac{1}{2}((1-\frac{\Delta\sigma}{\sigma})\frac{s-\mu}{\sigma})^2} ds \end{aligned}$$

We repeat the above process to eliminate the infinitesimal term $\Delta\sigma$ from the exponential term.

$$\begin{aligned} &\approx \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x \left(1 - \frac{\Delta\sigma}{\sigma}\right) \left(1 + \frac{(s-\mu)}{\sigma} \frac{\Delta\sigma}{\sigma}\right) \left(1 + \frac{(s-\mu)}{\sigma} \frac{\Delta\mu}{\sigma}\right) e^{-\frac{1}{2}(\frac{s-\mu}{\sigma})^2} ds \\ &\approx \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{1}{2}(\frac{s-\mu}{\sigma})^2} \left(1 - \frac{\Delta\sigma}{\sigma} + \frac{\Delta\sigma + \Delta\mu}{\sigma} \frac{(s-\mu)}{\sigma}\right) ds \\ &= (1 - \Delta\sigma/\sigma) G_\xi(x) - (\Delta\mu + \Delta\sigma)g_\xi(x). \end{aligned}$$

□

Then, we can rewrite the probability of being ranked in the j 'th place after a small deviation.

$$\begin{aligned} p_j(\xi', \xi) &= G_{\xi'}\left(q_\xi\left(1 - \frac{j}{n}\right)\right) - G_{\xi'}\left(q_\xi\left(1 - \frac{j+1}{n}\right)\right) \\ &\approx (1 - \Delta\sigma/\sigma) \frac{1}{n} + (\Delta\mu + \Delta\sigma) \left(g_\xi\left(q_\xi\left(1 - \frac{j+1}{n}\right)\right) - g_\xi\left(q_\xi\left(1 - \frac{j}{n}\right)\right)\right) \end{aligned} \quad (7)$$

By assumption 3.3, $(\Delta\mu + \Delta\sigma)$ is positive. Then, it's sufficient to show $g_\xi(q_\xi(1 - \frac{j+1}{n})) - g_\xi(q_\xi(1 - \frac{j}{n}))$ is decreasing in j . To make our life easier, we consider this in the continuous scale. Let $p = 1 - \frac{j+1}{n}$ and $\Delta p = \frac{1}{n}$. Then, let $f(p) = g_\xi(q_\xi(p)) - g_\xi(q_\xi(p + \Delta p))$ with $p \in (0, 1)$. We want to show that $f(p)$ is increasing in p .

First note that $\int_{-\infty}^{q_\xi(p)} g_\xi(x) dx = p$. Taking the derivative of p of both sides, we have $g_\xi(q_\xi(p)) = q'_\xi(p)^{-1}$. Thus, we want to show that $f(p) = q'_\xi(p)^{-1} - q'_\xi(p + \Delta p)^{-1}$ is increasing in p .

It is well known that the quantile of the Gaussian distribution can be represented by the inverse error function, i.e. $q(p) = \sqrt{2}\sigma \cdot \text{erf}^{-1}(2p-1) + \mu$ for a Gaussian with mean μ and std σ , where erf^{-1} is the inverse error function. Furthermore, we know the derivative of the inverse error function is $\frac{d}{dx} \text{erf}^{-1}(x) = \frac{1}{2}\sqrt{\pi}e^{(\text{erf}^{-1}(x))^2}$. Combining these,

$$\begin{aligned} \frac{\partial}{\partial p} f(p) &= \frac{1}{\sqrt{2\pi}\sigma} \cdot \frac{\partial}{\partial p} \left(e^{-(\text{erf}^{-1}(2p-1))^2} - e^{-(\text{erf}^{-1}(2(p+\Delta p)-1))^2} \right) \\ &= \frac{\sqrt{2}}{\sigma} (-\text{erf}^{-1}(2p-1) + \text{erf}^{-1}(2(p+\Delta p)-1)) \end{aligned}$$

Because $\text{erf}^{-1}(x)$ is increasing in x , we know $\frac{\partial}{\partial p} f(p)$ is positive which completes the proof. \square

A.2 The Optimal RO-Payment Function For Neutral Agents

Proof of Proposition 4.3. We start with solving the principal's optimization problem 4. Given that the agents are neutral we can write down the Lagrange and the KKT conditions as:

$$L(\hat{\mathbf{t}}, \alpha, \beta, \gamma) = \sum_j \hat{t}_j - \sum_j \alpha_j \hat{t}_j + \beta c(\xi) - \frac{\beta}{n} \sum_{j=1}^n \hat{t}_j - \gamma \sum_{j=1}^n p'_j(\xi) \cdot \hat{t}_j + \gamma c'(\xi).$$

- ① $\alpha_j = 1 - \frac{\beta}{n} - \gamma \cdot p'_j(\xi)$ for any $j \in [n]$;
- ② $\alpha_j \hat{t}_j = 0$ for any $j \in [n]$;
- ③ $\beta \cdot \left(c(\xi) - \frac{1}{n} \sum_{j=1}^n \hat{t}_j \right) = 0$;
- ④ $\sum_{j=1}^n p'_j(\xi) \cdot \hat{t}_j = c'(\xi)$;
- ⑤ $\alpha, \beta \geq 0$;
- ⑥ $-\hat{\mathbf{t}}, \left(c(\xi) - \frac{1}{n} \sum_{j=1}^n \hat{t}_j \right) \leq 0$.

Let $\omega(\xi) = c'(\xi)/p'_1(\xi)$. Now, we show that if IR is not binding, the solution to this problem is $\hat{t}_1 = \omega(\xi)$ and $\hat{t}_j = 0$ for any $j > 1$. IR is not binding implies $\beta = 0$ (condition ③). Then, we look at condition ①. Note that $\alpha_j \geq 0$ for any j and at least one of the α_j is equal to zero. Otherwise $\hat{t}_j = 0$ for any j (condition ②), and condition ④ is violated. There are two possible cases: if $\gamma < 0$, $\alpha_j = 0$ if and only if $p'_j(\xi)$ reaches its minimum; If $\gamma > 0$, $\alpha_i = 0$ if and only if $p'_j(\xi)$ reaches its maximum. (Note that $\gamma = 0$ is trivially infeasible.)

In lemma 4.2, we show that $p'_j(\xi)$ is decreasing in j given a fixed ξ . This property implies that the first case, i.e. $\gamma < 0$, is not feasible. Because $p'_j(\xi)$ reaches its minimum when $j = n$. However, if $\alpha_n = 0$ and $\hat{t}_n > 0$, condition ④ is violated given that c is increasing (RHS of ④ is positive) and $p'_j(\xi) < 0$ (LHS of ④ is negative). Therefore, the only possible solution is $\alpha_1 = 0$ and $\hat{t}_1 > 0$. By condition ④, $\hat{t}_1 = \omega(\xi)$ as $\Delta e \rightarrow 0^+$.

The above argument assumes IR is not binding, when is true when $\omega(\xi) \geq n \cdot c(\xi)$ or equivalently, $\eta(\xi) \geq n \cdot p'_j(\xi)$. If $\eta(\xi) < n \cdot p'_i(\xi)$, IR is binding, which implies that $\sum_{j=1}^n \hat{t}_j = n \cdot c(\xi)$. Any RO-payment function that satisfies FOC and makes IR binding are optimal. If we apply a threshold RO-payment function that pays agent j $\hat{t}_j = \tau$ if $1 \leq j \leq \hat{n}$, we completes the proof by solving for \hat{n} and τ . \square

A.3 The Optimal RO-Payment Function For Loss-averse Agents

The precise version of Proposition 4.4 is shown here. For simplification, let $\eta(\xi) = \frac{c'(\xi)}{c(\xi)}$. Then, let

$$H(\xi, k) = \left(\left(1 + \rho \frac{n-k}{n} \right) \eta(\xi) - \sum_{j=2}^k p'_j(\xi) + \rho \sum_{j=k+1}^n p'_j(\xi) \right),$$

and let $L(k) = (1 + \rho)(n - k) + 1$, we have the following results.

Proposition A.2. *Suppose \mathbf{p}' has rank-order impact at $\xi \in [0, 1]$, and agents are loss-averse.*

1. **IR is not binding:** *If $H(\xi, \bar{n}) \geq L(\bar{n}) \cdot p'_1(\xi)$, the optimal RO-payment function satisfies $\hat{t}_1 = H(\xi, \bar{n}) \cdot c(\xi)/p'_1(\xi)$, $\hat{t}_j = c(\xi)$ for $1 < j \leq \bar{n}$ and $\hat{t}_j = 0$ for $\bar{n} < j \leq n$, with threshold \bar{n} such that $(1 + \rho)p'_{\bar{n}}(\xi) = p'_1(\xi)$;*
2. **IR is binding:** *Otherwise, the optimal RO-payment function makes IR binding and pays fewer agents 0, where $\hat{t}_1 = H(\xi, \hat{n}) \cdot c(\xi)/p'_1(\xi)$, $\hat{t}_j = c(\xi)$ for $1 < j \leq \hat{n}$ and $\hat{t}_j = 0$ for $\hat{n} < j \leq n$, with threshold \hat{n} such that $H(\xi, \hat{n}) = L(\hat{n}) \cdot p'_1(\xi)$.*

Proof of Proposition 4.4 & A.2. Given the indifferentiability of the loss-averse utility, instead of using KKT conditions, we provide a more intuitive proof. As usual, we first ignore the IR constraint. Then the goal of the principal is to satisfied FOC with the minimum payments. Thus, starting with the all-zero payment, he will pay agents with the largest marginal return until FOC is satisfied. The marginal return of paying an agent with ranking j is

$$dt_j \begin{cases} = (1 + \rho)p'_j(\xi) & \text{if } \hat{t}_j < c(\xi), \\ \in [p'_j(\xi), (1 + \rho)p'_j(\xi)] & \text{if } \hat{t}_j = c(\xi), \\ = p'_j(\xi) & \text{if } \hat{t}_j > c(\xi). \end{cases}$$

Then, by Lemma 4.2, the optimal RO-payment function pays each agent j their cost $c(\xi)$ in the order of their ranking until some \bar{n} such that the principal is marginally better off to pay the top one agent more than $c(\xi)$ rather than paying the $\bar{n} + 1$ agent anything positive. The threshold \bar{n} therefore satisfies $(1 + \rho)p'_{\bar{n}}(\xi) = p'_1(\xi)$. Thus, the optimal RO-payment function is $\hat{t}_j = c(\xi)$ for $1 < j \leq \bar{n}$, $\hat{t}_j = 0$ for $\bar{n} < j \leq n$ and \hat{t}_1 such that FOC is satisfied. This gives us

$$\hat{t}_1 = \left(\left(1 + \rho \frac{n-k}{n} \right) c'(\xi) - c(\xi) \sum_{j=2}^k p'_j(\xi) + \rho c(\xi) \sum_{j=k+1}^n p'_j(\xi) \right) / p'_1(\xi).$$

Let $H(\xi, k) = \left(\left(1 + \rho \frac{n-k}{n} \right) \eta(\xi) - \sum_{j=2}^k p'_j(\xi) - \rho \sum_{j=k+1}^n p'_j(\xi) \right)$, then $\hat{t}_1 = H(\xi, \bar{n}) \cdot \frac{c(\xi)}{p'_1(\xi)}$. The condition for this to be true relies on IR being satisfied, i.e. $\hat{t}_1 + (\bar{n} - 1)c(\xi) \geq nc(\xi) + \rho(n - \bar{n})c(\xi)$. Let $L(k) = (1 + \rho)(n - k) + 1$. The condition becomes $H(\xi, \bar{n}) \geq L(\bar{n}) \cdot p'_1(\xi)$.

When IR is binding, i.e. $H(\xi, \bar{n}) < L(\bar{n}) \cdot p'_1(\xi)$, the payments satisfy $\sum_{j=1}^n \hat{t}_j = nc(\xi) + \sum_{j=1}^n \rho(c(\xi) - \hat{t}_j)^+$. Then, the goal is to minimize $\sum_{j=1}^n \rho(c(\xi) - \hat{t}_j)^+$, i.e. to overcome as more agents' cost as possible. With the same argument, the optimal ORPF pays agents with ranking smaller than some threshold \hat{n} their cost and pay the top one agent \hat{t}_1 such that FOC is satisfied and IR is binding. This gives us $\hat{t}_1 = H(\xi, \hat{n}) \cdot \frac{c(\xi)}{p'_1(\xi)}$ and \hat{n} such that $H(\xi, \hat{n}) = L(\hat{n}) \cdot p'_1(\xi)$. Note that $\hat{n} < n$ because when $\hat{n} = n$, IR is satisfied because everyone is paid her cost but FOC can never be satisfied because there is no incentive to exert higher effort.

Finally, we complete the proof by showing $\hat{n} \geq \bar{n}$. Note that $H(\xi, \bar{n}) < L(\bar{n}) \cdot p'_1(\xi)$ but $H(\xi, \hat{n}) = L(\hat{n}) \cdot p'_1(\xi)$. We only have to show that the marginal return of increasing k is positive for function $H(\xi, k) - L(k) \cdot p'_1(\xi)$. We have that the marginal return is $(1 + \rho)p'_1(\xi) - (1 + \rho)p'_k(\xi) \geq 0$, which completes the proof. \square

A.3.1 Inclusiveness Increases with Loss-aversion

Proof of Corollary 4.6. The proof is straightforward. With Proposition A.2, when IR is not binding, $n^I = \bar{n}$ which is determined by $(1 + \rho)p'_n(\xi) = p'_1(\xi)$. Because $p'_j(\xi)$ is decreasing in j by Lemma 4.2, \bar{n} is increasing in ρ .

When IR is binding, $n^I = \hat{n}$ is determined by $H(\xi, \hat{n}) = L(\hat{n}) \cdot p'_1(\xi)$. If we can show that $H(\xi, k) - L(k) \cdot p'_1(\xi)$ is decreasing in ρ , we can complete the proof because we know that $H(\xi, k) - L(k) \cdot p'_1(\xi)$ is increasing in k . It turns out the derivative of this term w.r.t. ρ is $\frac{n-k}{n}\eta(\xi) - \sum_{j=k+1}^n p'_j(\xi) - (n-k)p'_1(\xi)$. Because $\sum_{j=k+1}^n p'_j(\xi) \leq 0$ for any k , and $\frac{\eta(\xi)}{n} < p'_1(\xi)$ when $n \rightarrow \infty$, the derivative is negative and we complete the proof. \square

A.4 The Optimal RO-Payment Function For Risk-averse Agents

For risk-averse agents, $u_a(t_i, e_i) = r_a(t_i) - c(e_i)$. Let $\phi(x) = r_a^{-1}(x)$ be the inverse of the reward function, and ϕ' be the derivative. Let $v(j, k, \beta, \xi) = (\phi')^{-1} \left(\left(\phi'(0) - \frac{\beta}{n} \right) \cdot \frac{p'_j(\xi)}{p'_{k+1}(\xi)} + \frac{\beta}{n} \right)$.

Proposition A.3. *Suppose p' has rank-order impact at $\xi \in [0, 1]$, and agents are risk-averse.*

1. **IR is not binding:** If $\sum_{j=1}^{\bar{n}} v(j, \bar{n}, 0, \xi) \geq n \cdot c(\xi)$, the optimal RO-payment function satisfies $r_a(\hat{t}_j) = v(j, \bar{n}, 0, \xi)$ for $1 \leq j \leq \bar{n}$ and $\hat{t}_j = 0$ otherwise, with $\bar{n} \leq \frac{n}{2}$ determined by the FOC constraint, i.e. $\sum_{j=1}^{\bar{n}} p'_j(\xi) \cdot v(j, \bar{n}, 0, \xi) = c'(\xi)$;
2. **IR is binding:** Otherwise, the optimal RO-payment function satisfies $r_a(\hat{t}_j) = v(j, \hat{n}, \beta, \xi)$ for $1 \leq j \leq \hat{n}$ and $\hat{t}_j = 0$ otherwise, with $\hat{n} \geq \bar{n}$ and β determined by the FOC and IR constraints.

Proof of A.3. Because $\phi(x) = r_a^{-1}(x)$ is a differentiable convex function, the problem is a convex optimization problem. We can rewrite the principal's problem in terms of $r_j = r_a(\hat{t}_j)$ and write down the Lagrange and the KKT conditions.

$$L(\mathbf{r}, \alpha, \beta, \gamma) = \sum_j^n \phi(r_j) - \sum_j^n \alpha_j r_j + \beta c(\xi) - \frac{\beta}{n} \sum_{j=1}^n r_j - \gamma \sum_{j=1}^n p'_j(\xi) \cdot r_j + \gamma c'(\xi).$$

- ① $\alpha_j = \phi'(r_j) - \frac{\beta}{n} - \gamma \cdot p'_j(\xi)$ for any $j \in [n]$;
- ② $\alpha_j r_j = 0$ for any $j \in [n]$;
- ③ $\beta \cdot \left(c(\xi) - \frac{1}{n} \sum_{j=1}^n r_j \right) = 0$;
- ④ $\sum_{j=1}^n p'_j(\xi) \cdot r_j = c'(\xi)$;
- ⑤ $\alpha, \beta \geq 0$;
- ⑥ $-\mathbf{r}, \left(c(\xi) - \frac{1}{n} \sum_{j=1}^n r_j \right) \leq 0$.

Again, we start with the case where IR is not binding and $\beta = 0$. Thus, by ①, $\alpha_j = \phi'(r_j) - \gamma \cdot p'_j(\xi)$. Whenever $\hat{t}_j > 0$, $r_j > 0$ and $\alpha_j = 0$. By Lemma 4.2, $p'_j(\xi)$ is decreasing in j , and for the same reason in Appendix A.2, $\gamma > 0$. Therefore, the optimal payment scheme takes a threshold form for some threshold \bar{n} where $\hat{t}_j > 0$ for $1 \leq j \leq \bar{n}$ and $\hat{t}_j = 0$ otherwise. Furthermore, the payments satisfy that $\frac{\phi'(r_1)}{p'_1(\xi)} = \frac{\phi'(r_2)}{p'_2(\xi)} = \dots = \frac{\phi'(0)}{p'_{\bar{n}+1}(\xi)}$, or alternatively $r_j = (\phi')^{-1} \left(\phi'(0) \cdot \frac{p'_j(\xi)}{p'_{\bar{n}+1}(\xi)} \right)$. Note that because $r'_a(0) < \infty$, $\phi'(0) > 0$ and the solution is feasible. Then, to find the threshold \bar{n} , we can simply solve the FOC constraint, i.e. $\sum_{j=1}^{\bar{n}} p'_j(\xi) \cdot r_j = c'(\xi)$. The solution does not take a clean closed-form, but we know that $\bar{n} \leq \frac{n}{2}$ because $p'_j(\xi) \leq 0$ when $j \geq \frac{n}{2}$ (Eq. (7)), in which case $\alpha_j > 0$ for sure.

When IR is binding and $\beta > 0$, the same arguments still hold and $r_j = (\phi')^{-1} \left(\left(\phi'(0) - \frac{\beta}{n} \right) \cdot \frac{p'_j(\xi)}{p'_{\hat{n}+1}(\xi)} \right)$. Again, by solving IR is binding and FOC is satisfied, we have solutions for β and \hat{n} . Furthermore, we know that while fixing any ξ , in the case where IR is considered, the threshold \hat{n} is no less than \bar{n} which is the threshold when IR is not considered. First, if $\phi'(0) - \frac{\beta}{n} < 0$, $p'_{\hat{n}+1}(\xi) < 0$ and $\hat{n} \geq \frac{n}{2} \geq \bar{n}$. Second, if $\phi'(0) - \frac{\beta}{n} \geq 0$, suppose $\bar{n} > \hat{n}$. Every r_j in the IR-binding case is smaller than the case where IR is not binding. Consequently, ④ is violated which implies that $\bar{n} \leq \hat{n}$. \square

A.4.1 Inclusiveness is Not Monotone With Risk-aversion

Now, we show that more risk-averse agents does not imply more an inclusive optimal RO-payment function.

Corollary A.4. *Suppose $n \rightarrow \infty$ and agents are risk-averse. Let r_{a1} and r_{a2} be two concave reward functions of agents such that $\frac{\phi'_1(x)}{\phi'_1(0)} > \frac{\phi'_2(x)}{\phi'_2(0)}$ for any $x > 0$, where ϕ'_1 and ϕ'_2 are the derivative of the inverse of r_{a1} and r_{a2} respectively. Then, if IR is not binding, the optimal RO-payment function when agents have reward function r_{a1} is more inclusive than the case of r_{a2} . However, if IR is not binding, both cases are possible.*

Proof of A.4. First, we show that if IR is not binding, the RO-payment function is more inclusive as $\frac{\phi'(x)}{\phi'(0)}$ becomes larger for any $x > 0$. By Proposition A.3, when IR is not binding, the optimal RO-payment function is determined by

$$\frac{\phi'(r_a(\hat{t}_j))}{\phi'(0)} = \frac{p'_j(\xi)}{p'_{\bar{n}+1}(\xi)}. \quad (8)$$

Suppose $\frac{\phi'_1(x)}{\phi'_1(0)} > \frac{\phi'_2(x)}{\phi'_2(0)}$ for any $x > 0$, but \hat{t}_1 is more exclusive than \hat{t}_2 , i.e. $\bar{n}_1 < \bar{n}_2$. Then, for any $j \leq \bar{n}$, $\frac{p'_j(\xi)}{p'_{\bar{n}_1+1}(\xi)} < \frac{p'_j(\xi)}{p'_{\bar{n}_2+1}(\xi)}$ due to Lemma 4.2. As a result, to satisfy eq. (8), $r_{a1}(\hat{t}_{1,j}) < r_{a2}(\hat{t}_{2,j})$ for any $j \leq \bar{n}_1 \leq \bar{n}_2$. However, one of the payments, \hat{t}_1 or \hat{t}_2 must violate IR, which implies $\sum_{j=1}^{\bar{n}} r_a(\hat{t}_j) = c(\xi)$, because $\sum_{j=1}^{\bar{n}_1} r_a(\hat{t}_{1,j}) < \sum_{j=1}^{\bar{n}_2} r_a(\hat{t}_{2,j})$. Therefore, \hat{t}_1 must be at least as inclusive as \hat{t}_2 .

Second, we show that this pattern does not generally hold when IR is binding. Now, the optimal RO-payment function must satisfy

$$\frac{\phi'(r_a(\hat{t}_j)) - \frac{\beta}{n}}{\phi'(0) - \frac{\beta}{n}} = \frac{p'_j(\xi)}{p'_{\hat{n}+1}(\xi)}, \quad (9)$$

with $\beta > 0$. On one hand, the optimal RO-payment function can be more exclusive as $\frac{\phi'(x)}{\phi'(0)}$ increasing. Again, suppose $\frac{\phi'_1(x)}{\phi'_1(0)} > \frac{\phi'_2(x)}{\phi'_2(0)}$ for any $x > 0$ and $\phi'_1(0) = \phi'_2(0)$. In this case,

$$\frac{\phi'_1(x) - \frac{\beta}{n}}{\phi'_1(0) - \frac{\beta}{n}} - \frac{\phi'_2(x) - \frac{\beta}{n}}{\phi'_2(0) - \frac{\beta}{n}} = \frac{\phi'_1(x) - \phi'_2(x)}{\phi'_1(0) - \frac{\beta}{n}} > 0.$$

This implies that if $\phi'_1(0) - \frac{\beta}{n} > 0$, the same arguments in the IR not binding case still hold and \hat{t}_1 must be at least as inclusive as \hat{t}_2 .

On the other hand, \hat{t}_1 can be more exclusive when $\frac{\phi'_1(x)}{\phi'_1(0)} > \frac{\phi'_2(x)}{\phi'_2(0)}$ for any $x > 0$. Consider the case where $\phi'_1(x) - \phi'_2(x) > \phi'_1(0) - \phi'_2(0)$, $0 < \phi'_2(0) < \phi'_1(0) < \frac{\beta}{n}$. In this case,

$$\begin{aligned} \frac{\phi'_1(x) - \frac{\beta}{n}}{\phi'_1(0) - \frac{\beta}{n}} - \frac{\phi'_2(x) - \frac{\beta}{n}}{\phi'_2(0) - \frac{\beta}{n}} &= \frac{\phi'_1(x)\phi'_2(0) - \phi'_2(x)\phi'_1(0) + \frac{\beta}{n} \cdot (\phi'_2(x) - \phi'_1(x) + \phi'_1(0) - \phi'_2(0))}{(\phi'_1(0) - \frac{\beta}{n}) \cdot (\phi'_2(0) - \frac{\beta}{n})} \\ &< \frac{\phi'_1(x)\phi'_2(0) - \phi'_2(x)\phi'_1(0) + \phi'_2(0) \cdot (\phi'_2(x) - \phi'_1(x) + \phi'_1(0) - \phi'_2(0))}{(\phi'_1(0) - \frac{\beta}{n}) \cdot (\phi'_2(0) - \frac{\beta}{n})} \\ &= \frac{(\phi'_1(0) - \phi'_2(0)) \cdot (\phi'_2(0) - \phi'_2(x))}{(\phi'_1(0) - \frac{\beta}{n}) \cdot (\phi'_2(0) - \frac{\beta}{n})} \\ &\leq 0. \end{aligned}$$

This implies that when agents become more risk-averse, i.e. $\frac{\phi'_1(x)}{\phi'_1(0)} > \frac{\phi'_2(x)}{\phi'_2(0)}$ for any $x > 0$, the LHS of eq. (9) becomes smaller. Now, suppose $\hat{n}_1 > \hat{n}_2$. We have $p'_{\hat{n}_1+1} < p'_{\hat{n}_2+1} < 0$, and thus $\frac{p'_j(\xi)}{p'_{\hat{n}_1+1}(\xi)} > \frac{p'_j(\xi)}{p'_{\hat{n}_2+1}(\xi)}$ for any $j \leq \hat{n}_2$. As a result, to satisfy eq. (9), $r_{a1}(\hat{t}_{1,j}) > r_{a2}(\hat{t}_{2,j})$ for any $j \leq \hat{n}_2$. Again, this violates the IR constraint for the same reason in the IR not binding case, which implies $\hat{n}_1 \leq \hat{n}_2$. \square

A.5 Sensitivity is A Sufficient Statistic

Proof of Proposition 4.8. While fixing ξ and ξ' , we view $p_j(\xi', \xi)$ as a function of $\mu(\xi)$ and $\sigma(\xi)$, denoted as $p_j(\mu, \sigma, \xi', \xi)$. We want to show that if a performance measurement has a higher sensitivity, it requires weakly lower payment to elicit a goal effort.

The intuition of the proof is that suppose \hat{t}^* is the optimal RO-payment function when performance measurement Ψ is applied. Now, under Assumption 4.1, fixing ξ , if $\delta(\xi)$ increases, we show that the FOC constraint is easier to be satisfied, i.e. FOC can be satisfied with strictly lower total payment. This implies that with a performance measurement that has higher sensitivity, the principal is at least not worse-off. To see this, when IR is not binding, the principal can reduce the payment to satisfy FOC without violating IR and LL if a more sensitive performance measurement is applied. When IR is binding, the principal can reduce \hat{t}_1 by ϵ_1 and increase \hat{t}_n by $\epsilon_n \leq \epsilon_1$ such that FOC is satisfied and IR is still binding. In particular, for neutral agents, $\epsilon_n = \epsilon_1$, in which case the principal is equivalent; for risk/loss-averse agents, $\epsilon_n \leq \epsilon_1$ because further increasing the payment to the top agent has a (weakly) smaller marginal return than using that payment to reward lower ranked agents, in which case the principal is (weakly) better-off.²³

With this intuition, our goal is to show that FOC can be satisfied with strictly lower payment as δ increases. Let $\lambda_j = r_a(\hat{t}_j) - \rho(c(\xi') - \hat{t}_j)^+$. Note that the FOC constraint requires that $\sum_{j=1}^n (p_j(\mu, \sigma, \xi', \xi) - \frac{1}{n}) \cdot \lambda_j = c'(\xi)$. Note that the only term that depends on $\delta(\xi)$ is $p_j(\mu, \sigma, \xi', \xi)$. Thus, the rest of the proof can be summarized in Lemma A.5, which shows that the left-hand-side of the FOC constraint is increasing in δ while fixing the payment \hat{t} , or equivalently, FOC can be satisfied with lower payment as δ increases.

We then complete the proof by showing $\lambda_j = r_a(\hat{t}_j) - \rho(c(\xi') - \hat{t}_j)^+$ is decreasing in j under the optimal RO-payment function for any type of agents. This can be proven by showing that the optimal RO-payment function is monotone decreasing with $\hat{t}_j \leq \hat{t}_k$ if $j \geq k$. According to Proposition 4.3, A.2 and A.3, this is exactly the case. \square

Lemma A.5. *Under Assumption 4.1, for any $\xi \in [0, 1]$, $\sum_{j=1}^n p_j(\mu, \sigma, \xi', \xi) \cdot \lambda_j$ is increasing in $\delta(\xi) = \frac{\mu'(\xi)}{\sigma(\xi)}$ if $0 < \lambda_j \leq \lambda_k$ for any $1 \leq j \leq n$.*

Proof. Fixing ξ , let μ, σ be the mean and std of the score distribution while all agents exert an effort ξ . Then, let agent i deviate to a slightly higher effort ξ' with $\xi' - \xi \rightarrow 0$. Let $\Delta\mu$ and $\Delta\sigma$ be the change of the mean and std of agent i 's score distribution. By Assumption 4.1, $\Delta\mu \gg \Delta\sigma$, and because $\mu(e)$ is differentiable,

²³For loss-averse agents, the argument is strict when the payment to agent n , \hat{t}_n , is strictly smaller than $c(\xi)$; for risk-averse agents, the argument is strict when \hat{t}_n is strictly smaller than \hat{t}_1 and r_a is strictly concave.

$\Delta\mu \rightarrow 0$. Then, with the same approach in Lemma A.1, we can rewrite the probability $p(\mu, \sigma, \xi', \xi, j)$ as

$$\begin{aligned}
& p_j(\mu, \sigma, \xi', \xi) \\
&= \int_{-\infty}^{\infty} g(\xi + \Delta e, x) \binom{n-1}{j-1} (G(\xi, x))^{n-j} (1 - G(\xi, x))^{j-1} dx \\
&= \frac{1}{\sqrt{2\pi}(\sigma + \Delta\sigma)} \int_{-\infty}^{\infty} e^{-\frac{1}{2}\left(\frac{x-\mu-\Delta\mu}{\sigma+\Delta\sigma}\right)^2} \binom{n-1}{j-1} \left(\frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{1}{2}\left(\frac{s-\mu}{\sigma}\right)^2} ds \right)^{n-j} \left(1 - \frac{1}{\sigma\sqrt{2\pi}} \int_{-\infty}^x e^{-\frac{1}{2}\left(\frac{s-\mu}{\sigma}\right)^2} ds \right)^{j-1} dx \\
&= \frac{1}{\sqrt{2\pi}(\sigma + \Delta\sigma)} \int_{-\infty}^{\infty} e^{-\frac{1}{2}\left(\frac{x-\mu-\Delta\mu}{\sigma+\Delta\sigma}\right)^2} \binom{n-1}{j-1} \left(\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{x-\mu}{\sigma}} e^{-\frac{1}{2}y^2} dy \right)^{n-j} \left(1 - \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\frac{x-\mu}{\sigma}} e^{-\frac{1}{2}y^2} dy \right)^{j-1} dx \\
&\hspace{25em} (y = \frac{s-\mu}{\sigma}) \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \left(1 - \frac{\Delta\sigma}{\sigma} \right) e^{-\frac{1}{2}\left(\left(1-\frac{\Delta\sigma}{\sigma}\right)z - \frac{\Delta\mu}{\sigma}\right)^2} \binom{n-1}{j-1} G_0(z)^{n-j} (1 - G_0(z))^{j-1} dz \\
&\hspace{25em} (z = \frac{x-\mu}{\sigma}) \\
&\approx \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \left(1 + \frac{\Delta\mu}{\sigma} z \right) e^{-\frac{1}{2}z^2} \binom{n-1}{j-1} G_0(z)^{n-j} (1 - G_0(z))^{j-1} dz,
\end{aligned}$$

where the last step is derived using the same recipe as Lemma A.1 and then omitting the terms of $\Delta\sigma$. Let $\delta = \frac{\Delta\mu}{\sigma}$. We have,

$$\begin{aligned}
& \sum_{j=1}^n \lambda_j \frac{\partial p_j}{\partial \delta} \\
&= \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} z e^{-\frac{1}{2}z^2} \sum_{j=1}^n \lambda_j \binom{n-1}{j-1} \left(G_0(z)^{n-j} (1 - G_0(z))^{j-1} \right) dz \\
&= \frac{1}{\sqrt{2\pi}} \int_0^{\infty} z e^{-\frac{1}{2}z^2} \sum_{j=1}^n \lambda_j \binom{n-1}{j-1} \left(G_0(z)^{n-j} (1 - G_0(z))^{j-1} \right) dz \\
&\quad - \frac{1}{\sqrt{2\pi}} \int_0^{\infty} z e^{-\frac{1}{2}z^2} \sum_{j=1}^n \lambda_j \binom{n-1}{j-1} \left((G_0(-z))^{n-j} (1 - G_0(-z))^{j-1} \right) dz
\end{aligned}$$

Then, we reorder the terms in the second summation such that every j in the second summation is replaced with $n - j - 1$, while the summation remains the same. In this way, every term in the first summation with weight λ_j is paired with a term in the second summation with weight λ_{n-j+1} . Furthermore, because $\binom{n-1}{j-1} = \binom{n-1}{n-j}$ and $G_0(-z) = 1 - G_0(z)$, this allows us to term-wisely combine the two summations in the above equation.

$$= \frac{1}{\sqrt{2\pi}} \int_0^{\infty} z e^{-\frac{1}{2}z^2} \sum_{j=1}^n (\lambda_j - \lambda_{n-j+1}) \binom{n-1}{j-1} \left(G_0(z)^{n-j} (1 - G_0(z))^{j-1} \right) dz$$

Next, within the integral, we first double the summation by adding another summation over k , where $k = n - j + 1$ takes the inverse order of the original summation, then we divide the sum by 2, which does not change the quantity.

$$\begin{aligned}
&= \frac{1}{\sqrt{2\pi}} \int_0^{\infty} z e^{-\frac{1}{2}z^2} \frac{1}{2} \left(\sum_{j=1}^n (\lambda_j - \lambda_{n-j+1}) \binom{n-1}{j-1} \left(G_0(z)^{n-j} (1 - G_0(z))^{j-1} \right) \right. \\
&\quad \left. + \sum_{k=1}^n (\lambda_{n-k+1} - \lambda_k) \binom{n-1}{n-k} \left(G_0(z)^{k-1} (1 - G_0(z))^{n-k} \right) \right) dz
\end{aligned}$$

Again, we and combine these two summations term by term, i.e. pairing $j = i$ with $k = i$ for every i from 1 to n .

$$= \frac{1}{2\sqrt{2\pi}} \int_0^\infty z e^{-\frac{1}{2}z^2} \left(\sum_{i=1}^n \binom{n-1}{i-1} (\lambda_i - \lambda_{n-i+1}) \left(G_0(z)^{n-i} (1 - G_0(z))^{i-1} - G_0(z)^{i-1} (1 - G_0(z))^{n-i} \right) \right) dz$$

$$\geq 0.$$

The integral is non-negative because when $z \geq 0$, $G_0(z) \geq 1 - G_0(z)$. Therefore, the term $G_0(z)^{n-i} (1 - G_0(z))^{i-1} - G_0(z)^{i-1} (1 - G_0(z))^{n-i}$ is non-negative for every $1 \leq i \leq \lfloor \frac{n+1}{2} \rfloor$, and negative otherwise. Furthermore, because λ_i is decreasing with i , the term $\lambda_i - \lambda_{n-i+1}$ is also non-negative for every $1 \leq i \leq \lfloor \frac{n+1}{2} \rfloor$, and negative otherwise. Thus, the product of these two terms is always non-negative, which implies the integrand is non-negative. This completes the proof. \square

B Proofs and Details of Section 5

B.1 Higher Mean and Higher Variance Benefit Untruthful Deviations

Proof of Lemma 5.1. For simplicity, we first normalize the score distributions such that all agents but i has a score follows the standard Gaussian distribution, while agent i 's score follows $g' = \mathcal{N}(\mu', \sigma')$. By assumption, $\mu' \leq 0$. Under the winner-take-all tournament with a fixed prize, agent i 's expected payment is proportional to the probability of being ranked first. We denote this probability as $p_1(\mu', \sigma')$.

To prove the first part of the lemma, we show that the first derivatives of p_1 with respect to both μ' and σ' are positive.

$$\begin{aligned} p_1(\mu, \sigma) &= \int_{-\infty}^{+\infty} g'(x) G_0(x)^{n-1} dx \\ &= \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi}\sigma} e^{-\frac{1}{2}\left(\frac{x-\mu'}{\sigma}\right)^2} G_0(x)^{n-1} dx \\ &= \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}y^2} G_0(\sigma'y + \mu')^{n-1} dy. \end{aligned}$$

The first-order derivatives can be written as

$$\begin{aligned} \frac{\partial p_1}{\partial \mu'} &= \frac{n-1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} e^{-\frac{1}{2}y^2} G_0(\sigma'y + \mu')^{n-2} g_0(\sigma'y + \mu') dy > 0. \\ \frac{\partial p_1}{\partial \sigma'} &= \frac{n-1}{\sqrt{2\pi}} \int_{-\infty}^{+\infty} y e^{-\frac{1}{2}y^2} G_0(\sigma'y + \mu')^{n-2} g_0(\sigma'y + \mu') dy \\ &= \frac{n-1}{\sqrt{2\pi}} \int_0^{+\infty} y e^{-\frac{1}{2}y^2} (g_0(\sigma'y + \mu') G_0(\sigma'y + \mu')^{n-2} - g_0(-\sigma'y + \mu') G_0(-\sigma'y + \mu')^{n-2}) dy \\ &\geq \frac{n-1}{\sqrt{2\pi}} \int_0^{+\infty} y e^{-\frac{1}{2}y^2} g_0(\sigma'y + \mu') ((G_0(\sigma'y + \mu')^{n-2} - G_0(-\sigma'y + \mu')^{n-2}) dy \\ &> 0. \end{aligned}$$

For the derivative over σ' , the first inequality holds because when $\mu' \leq 0$, by the symmetry of the p.d.f. of the Gaussian distribution, $g_0(-\sigma'y + \mu') \geq g_0(\sigma'y + \mu')$ for any $y \geq 0$. Therefore, replacing $g_0(-\sigma'y + \mu')$ with $g_0(\sigma'y + \mu')$ lower bounds the integral. Furthermore, because $G_0(\sigma'y + \mu') > G_0(-\sigma'y + \mu')$ for any $y > 0$, the integral is positive.

For the second part of the proof, we want to show that while fixing μ' , a sufficiently large σ' will lead to a larger p_1 than obtaining the same score distribution as all other agents. Note that if all agents' score distributions are the same, $p_1 = \frac{1}{n}$ by symmetry.

Because while fixing μ' , p_1 is monotone increasing in σ' , we only have to show that $p_1 > \frac{1}{n}$ when $\sigma' \rightarrow \infty$. This holds because when $n \geq 3$, suppose $\gamma \sim \mathcal{N}(\mu, \sigma)$ and suppose $x \geq \mu$ (the analysis of $x < \mu$ is analogue), then

$$\lim_{\sigma \rightarrow \infty} \Pr(\gamma > x) = \Pr(\gamma > \mu) + \lim_{\sigma \rightarrow \infty} \Pr(\gamma \in [\mu, x]) = \Pr(\gamma > \mu) = \frac{1}{2} > \frac{1}{n}.$$

This completes the proof. \square

B.2 Adding Noise Helps Truthfulness

Proof of Proposition 5.3. Let agent i deviate from the truth-telling strategy profile and play an arbitrary strategy $\theta \in \Theta$. For simplicity, we first normalize the score distributions such that all agents but i has a score follows $\mathcal{N}(0, \sigma_\tau)$. Then, we bound the probability of being ranked first, p_1 , by identifying the “worst” possible deviation with the largest p_1 . Recall that we assume the strategy space is compact and the mean and standard deviation domains of the score distributions of all strategy profile are compact (see Assumption 3.4). Suppose the maximum mean of a unilateral untruthful deviation while all other agents are truthfully reporting is $\tilde{\mu} = \max_{\theta \in \Theta \setminus \{\tau\}} \mu(\theta, \tau)$, and the maximum standard deviation is $\tilde{\sigma} = \max_{\theta \in \Theta \setminus \{\tau\}} \sigma(\theta, \tau)$. Because the performance measurement is strongly truthful (Definition 3.2) and any unilateral deviation changes the deviating agent’s expected score more than other agents’ expected score (Assumption 3.5), $\tilde{\mu} < 0$. Then, by Lemma 5.1, no strategy $\theta \in \Theta$ can bring a higher p_1 than the one that induces a score distribution of $\tilde{g} = \mathcal{N}(\tilde{\mu}, \tilde{\sigma})$. Note that \tilde{g} may not be induced by any $\theta \in \Theta$. Then, the proof follows by showing that after adding a sufficiently large noise, even in the worst case, truth-telling still leads to the largest p_1 .

There are two cases. First, if $\tilde{\sigma} \leq \sigma_\tau$, by Proposition 5.2, no untruthful unilateral deviation can outperform truthful-telling. Therefore, in the proof that follows, we consider the case where $\tilde{\sigma} > \sigma_\tau$.

Note that the sum of two Gaussian variables also follows the Gaussian distribution. Therefore, when agents are truthful, the modified performance score follows $g'_\tau = \mathcal{N}(0, \sigma'_\tau)$ where $\sigma'_\tau = \sqrt{\sigma_\tau^2 + \sigma_\epsilon^2}$. We denote the c.d.f. of this distribution as G'_τ . Furthermore, for the worst possible deviation, the modified performance score follows $\tilde{g}' = \mathcal{N}(\tilde{\mu}, \tilde{\sigma}')$ where $\tilde{\sigma}' = \sqrt{\tilde{\sigma}^2 + \sigma_\epsilon^2}$. Then, we rewrite the probability of winning the first prize into the integral of standard Gaussian distributions.

$$\begin{aligned} p_1 &= \int_{-\infty}^{+\infty} \tilde{g}'(x) G'_\tau(x)^{n-1} dx \\ &= \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi}\tilde{\sigma}'} e^{-\frac{1}{2}\left(\frac{x-\tilde{\mu}}{\tilde{\sigma}'}\right)^2} \left(\int_{-\infty}^x \frac{1}{\sqrt{2\pi}\sigma'_\tau} e^{-\frac{1}{2}\left(\frac{y}{\sigma'_\tau}\right)^2} dy \right)^{n-1} dx \\ &= \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi}\tilde{\sigma}'} e^{-\frac{1}{2}\left(\frac{x-\tilde{\mu}}{\tilde{\sigma}'}\right)^2} \left(\int_{-\infty}^{\frac{x}{\sigma'_\tau}} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2} dt \right)^{n-1} dx \quad (\text{Let } t = \frac{y}{\sigma'_\tau}) \\ &= \int_{-\infty}^{+\infty} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}z^2} \left(\int_{-\infty}^{\frac{\tilde{\sigma}'}{\sigma'_\tau}z + \frac{\tilde{\mu}}{\sigma'_\tau}} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}t^2} dt \right)^{n-1} dz \quad (\text{Let } z = \frac{x-\tilde{\mu}}{\tilde{\sigma}'}) \\ &= \int_{-\infty}^{+\infty} g_0(z) G_0\left(\frac{\tilde{\sigma}'}{\sigma'_\tau}z + \frac{\tilde{\mu}}{\sigma'_\tau}\right)^{n-1} dz, \end{aligned}$$

where g_0 and G_0 are the p.d.f. and c.d.f. of the standard Gaussian. We want to show that if σ_ϵ is large enough, p_1 is smaller than $\frac{1}{n}$, the probability of winning the first prize while being truthful in the symmetric equilibrium.

$$\begin{aligned} p_1 - \frac{1}{n} &= \int_{-\infty}^{+\infty} g_0(x) \left(G_0\left(\frac{\tilde{\sigma}'}{\sigma'_\tau}x + \frac{\tilde{\mu}}{\sigma'_\tau}\right)^{n-1} - G_0(x)^{n-1} \right) dx \\ &= \int_{-\infty}^{+\infty} g_0\left(z - \frac{\tilde{\mu}}{\tilde{\sigma}' - \sigma'_\tau}\right) \left(G_0\left(\frac{\tilde{\sigma}'}{\sigma'_\tau}z - \frac{\tilde{\mu}}{\tilde{\sigma}' - \sigma'_\tau}\right)^{n-1} - G_0\left(z - \frac{\tilde{\mu}}{\tilde{\sigma}' - \sigma'_\tau}\right)^{n-1} \right) dz. \quad (z = x + \frac{\tilde{\mu}}{\tilde{\sigma}' - \sigma'_\tau}) \end{aligned}$$

This step helps us compare the point-wise value of two Gaussians, by creating a common bias, $\frac{\tilde{\mu}}{\tilde{\sigma}' - \sigma'_\tau}$. Next, because $\frac{1}{\tilde{\sigma}' - \sigma'_\tau} = \frac{1}{\sqrt{\tilde{\sigma}^2 + \sigma_\epsilon^2} - \sqrt{\sigma_\tau^2 + \sigma_\epsilon^2}} = \frac{\sqrt{\tilde{\sigma}^2 + \sigma_\epsilon^2} + \sqrt{\sigma_\tau^2 + \sigma_\epsilon^2}}{\tilde{\sigma}^2 - \sigma_\tau^2} \approx \frac{2\sigma_\epsilon}{\tilde{\sigma}^2 - \sigma_\tau^2}$,

$$\approx \int_{-\infty}^{+\infty} g_0\left(z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) \left(G_0\left(\frac{\tilde{\sigma}'}{\sigma'_\tau} z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} - G_0\left(z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} \right) dz. \quad (\sigma_\epsilon \gg \tilde{\sigma})$$

Then, we break the integral into two parts, i.e. the integral from $-\infty$ to 0 and the integral from 0 to ∞ . For the integral over the negative space, we replace the variable z with $y = -z$. With the symmetry of the standard Gaussian, i.e. $g_0(-x) = g_0(x)$, we can rewrite the above equations

$$\begin{aligned} &= \int_0^{+\infty} g_0\left(z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) \left(G_0\left(\frac{\tilde{\sigma}'}{\sigma'_\tau} z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} - G_0\left(z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} \right) dz \\ &+ \int_0^{+\infty} g_0\left(y + \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) \left(G_0\left(-\frac{\tilde{\sigma}'}{\sigma'_\tau} y - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} - G_0\left(-y - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} \right) dy. \end{aligned}$$

Next, because σ_ϵ is sufficiently large and $\tilde{\mu} < 0$, $g_0\left(z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) \approx 0$ when $z > 0$. Therefore, the second term in the above equation dictates the summation. Furthermore, because $\tilde{\sigma}' > \sigma'_\tau$ and $\tilde{\mu} < 0$, $G_0\left(-\frac{\tilde{\sigma}'}{\sigma'_\tau} z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) < G_0\left(-z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)$ for any $z > 0$. We have

$$\begin{aligned} &\approx \int_0^{+\infty} g_0\left(z + \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right) \left(G_0\left(-\frac{\tilde{\sigma}'}{\sigma'_\tau} z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} - G_0\left(-z - \frac{2\sigma_\epsilon \tilde{\mu}}{\tilde{\sigma}^2 - \sigma_\tau^2}\right)^{n-1} \right) dz \\ &< 0. \end{aligned}$$

Therefore, when the error distribution is diffuse enough, any untruthful deviation will never be preferred by the agent under the winner-take-all tournament, which completes the proof. \square

C The Fairness-Seeking Principal

Here, we provide a variance of our standard principal model. We have been focusing on the risk-neutral principal who aims to minimize the total payment given a fixed goal effort level. However, the principal may want to pay the agents with surplus to trade-off the efficiency and the fairness of the payments for various reasons. For example, the principal wants to reduce the variance of the payments even though he has to pay more due to an intrinsic notion of fairness, social pressure, or the low participation rate. We model the fairness-seeking principal with a penalty term in their utility function, $\Theta(\hat{\mathbf{t}})$. In this way, the principal aims to solve the same problem in (4) but to minimize the linear combination of the total payment and the fairness cost, i.e.

$$\min_{\hat{\mathbf{t}}} \sum_{j=1}^n \hat{t}_j + \lambda \cdot \Theta(\hat{\mathbf{t}}).$$

To simplify the analysis and provide intuitions, we consider two examples of the penalty functions while assuming the agents are neutral. First, similar to the loss-aversion case, let $\Theta(\hat{\mathbf{t}}) = \sum_j (c - \hat{t}_j)^+$ where c is a positive constant, e.g. $c = c(\xi)$ with a goal effort ξ . This example models the fact that the principal wants to pay the agents some money to overcome their cost of effort. With the same arguments as in Appendix A.3, one can easily verify that the optimal RO-payment function in this case is similar to the loss-averse agents case as shown in Fig. 2. In the optimal RO-payment function, the principal pays a fraction of agents who are ranked higher than some threshold c and the top one agent more than c . The only difference lies in the threshold. For example, when IR is not binding, the optimal threshold \bar{n} satisfies $p'_{\bar{n}+1}(\xi) = (1 - \lambda)p'_1(\xi)$ instead of $(1 + \rho)p'_{\bar{n}+1}(\xi) = p'_1(\xi)$ in Proposition 4.4.

For another example, let $\Theta(\hat{\mathbf{t}}) = \frac{1}{n} \sum_{j=1}^n (\hat{t}_j - \bar{t})^2$, where $\bar{t} = \frac{1}{n} \sum_{j=1}^n \hat{t}_j$. In this case, the principal aims to reduce the variance of the payments. Since the LL, IR and FOC constraints are linear in $\hat{\mathbf{t}}$, the principal's problem is convex which can be numerically solved. We thus use our ABM to learn the probability $p'_j(\xi)$ and then explore the trade-off between the efficiency, i.e. the total payment, and the fairness, i.e. the variance of the payments.




Figure 9: The trade-off between the efficiency and the fairness of the payments with λ varying from 0 to 1. In this example, we learn $p'_j(\xi)$ from the dataset of $W1$ with performance measurement the SC-Acc and $\xi = 0.8$. The cost function is $c(x) = x^8$.

In Fig. 9, we visualize the principal's trade-off between lowering the variance and saving the budget. The principal can trade-off the fairness and the efficiency by implementing a more inclusive mechanism which pays the agents more money than their cost of effort and use the surplus to reduce the variance. The trade-off depends on the setting, i.e. the cost function, goal effort and the confusion mapping. However, our numerical results suggest that the principal can greatly lower the variance without too much more cost of budget. For example, the variance can be reduced by 50% with only a 7% increase in the total payment in the example of Fig. 9.