

Report-Sensitive Spot-Checking in Peer-Grading Systems

Extended Abstract

Hedayat Zarkoob
University of British Columbia
Vancouver, BC
hzarkoob@cs.ubc.ca

Hu Fu
University of British Columbia
Vancouver, BC
hufu@cs.ubc.ca

Kevin Leyton-Brown
University of British Columbia
Vancouver, BC
kevinlb@cs.ubc.ca

KEYWORDS

Peer Grading, Peer Prediction, Mechanism Design

ACM Reference Format:

Hedayat Zarkoob, Hu Fu, and Kevin Leyton-Brown. 2019. Report-Sensitive Spot-Checking in Peer-Grading Systems. In *Proc. of the 18th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, Montreal, Canada, May 13–17, 2019, IFAAMAS, 3 pages.

1 INTRODUCTION

Peer grading systems make large courses more scalable, provide students with faster and more detailed feedback, and help teach students to think critically about the work of others. Various recent implementations of peer grading mechanisms make such systems relatively easy to deploy in practice [2, 11, 24]. The broader adoption of such systems faces a common, critical obstacle: motivating students to provide accurate grades. A natural solution is asking multiple students to grade the same assignment and rewarding them based on their behavior (e.g., based on the extent to which their grades agree with the grades given by other students). Such solutions have been explored in detail in a large literature on *peer prediction*, which considers how to incentivize agents to truthfully disclose unverifiable private information [4, 7–10, 12–17, 22, 23]. Unfortunately, almost all known peer prediction mechanisms also give rise to uninformative equilibria in which agents do not reveal their private information; e.g., all students grading an assignment favorably regardless of its quality [1, 8, 10, 17, 22]. Human experiments show that such strategic behavior does arise in practice [5].

Much subsequent work has attempted to identify peer prediction mechanisms in which the truthful equilibrium is always preferred by agents, or even in which no uninformative equilibria exist [1, 3, 8, 10, 17, 22]. One drawback of all such approaches is that they cannot do better than Nash equilibrium implementations. This is because agents' payoffs depend on other agents' actions, and so agents must reason about each other's behavior. In a classroom setting, where some students will almost surely fail to invest effort, students may need stronger incentives; we thus seek dominant strategy mechanisms. Such mechanisms can be obtained by incorporating trusted graders (TAs) more fundamentally into the mechanism: guaranteeing that each student report is, with some sufficiently large probability, compared to a trusted evaluation (which we will call a *spot check*), rather than to other student evaluations. The idea of combining such "spot checking" with peer grading mechanisms to incentivize accurate grading has been explored in

some recent past work [6, 7, 21, 24]. Because spot checking is expensive (e.g., TAs need to be paid in proportion to the amount of work they do), it is natural to seek to minimize the amount of spot checking required to obtain dominant strategies. This minimization problem was first attacked by Gao et al. [6], who proposed a very simple mechanism that makes truthfulness a dominant strategy by unconditionally rewarding students when they are not spot checked and penalizing them to the extent that they disagree with the TA otherwise. They compared this mechanism with various alternatives based on peer prediction, showing that the latter require strictly more spot checking than the former, even despite the fact that peer-prediction-based mechanisms do not offer dominant strategies.

Gao et al.'s model always performs spot checks with some fixed probability. It is intuitive to think that *report-sensitive spot checking*—that is, varying the spot checking probability based on the students' reports—could lower the expected amount of spot checking required overall. For example, imagine that an instructor already knows that a given problem set is extremely difficult. If the reported grades for a given submission are all very high, the instructor might believe that there is an increased likelihood that students have reported dishonestly, and so might want to spot check with a higher probability.

For the first time, this paper identifies the optimal dominant-strategy incentive-compatible report-sensitive spot-checking mechanism, which requires less spot checking than the previous state of the art, the simple mechanism of Gao et al. Like much other work in the literature [e.g., 1, 21], our analysis is limited to the case where students are asked to report only positive or negative grades about each assignment. Our new mechanism is general in several important senses: it allows for arbitrary numbers of graders per assignment and nearly arbitrary¹ prior probability distributions over both these signals and the noise models describing the probabilities that students and TAs will observe each signal given the ground truth.

One final and very recent related paper is worth mentioning here. Wang et al. [21] proposed a different approach for designing peer grading systems that also varies spot check probabilities. Their model, which assumes that TAs can directly observe whether a student invested effort, is substantially different from our own, and hence their mechanism is not directly applicable to our setting.

2 MODEL

A single assignment needs to be graded by a set N of students (with $|N| = n$) and has an unobservable binary quality $q \in Q = \{a, b\}$

¹The only assumption we make on these distributions is that spot checking each student with probability one yields dominant strategies in the mechanism of Gao et al.

drawn from a commonly known distribution $\Pr[q]$. Each student i , by exerting effort at cost c , can examine the submission and observe a signal $s_i \in Q$ that is informative about the assignment's quality. More formally, in a way that depends on the true quality q , the signals observed by different students are independently drawn from a commonly known distribution $\Pr[s|q]$.

In addition to the students, a teaching assistant (TA) may also receive a signal s_{TA} . Formally, s_{TA} is drawn from $\Pr[s|q]$; conditioning on the quality q , it is independent from the students' signals.

Strategy space. In our model, each student faces two strategic choices: whether to expend effort grading the assignment and what grade to report. Three actions are thus possible: the student (i) may be *truthful*, investing effort to examine the assignment, observing her signal, and reporting this signal; (ii) may invest effort but report a different signal than the one she observed; or (iii) may choose not to invest effort and report an arbitrary signal. In contrast, the TA is not a strategic agent. When asked to grade the assignment, the TA always reports an independently observed signal.

3 SPOT-CHECKING MECHANISMS

A spot-checking mechanism takes in students' reported signals and decides both whether a TA signal is needed and how much to reward the students.

Definition 3.1 (Spot-checking mechanism). A spot-checking mechanism is defined by a tuple (x_a, x_b, Y) , where:

- (1) $x_a : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$ defines the probability of spot checking an agent who reports a . Given two natural numbers (k, n) specifying the number of a 's reported by the agents and the total number of agents, x_a returns the probability that the mechanism will spot check agents reporting a .
- (2) $x_b : \mathbb{N} \times \mathbb{N} \rightarrow [0, 1]$ is an analogous function for computing the probability of spot checking an agent who reports b .
- (3) $Y : Q \times Q \rightarrow \mathbb{R}^+$ denotes the reward given to a student who is spot checked. $Y(r, s_{TA})$ is the reward given to a spot-checked student who made report r when the TA reported signal s_{TA} . When a student is not spot checked, she receives no reward.

Throughout this paper we focus on mechanisms where the reward function Y is the output agreement reward function, i.e., $Y(r, s_{TA}) = 1$ if $r = s_{TA}$, and 0, otherwise. This function has been widely studied in the peer prediction literature [18–20, 22].

Definition 3.2 (DSIC). A spot-checking mechanism is *dominant strategy incentive compatible* (DSIC) if, for each student i and for any strategies that the other students choose, i 's utility-maximizing strategy is to be truthful, i.e., to invest effort to observe her signal and to report what she observes.

3.1 ROS, RSS, and RSUS Mechanisms

Definition 3.3 (ROS Mechanism [6]). A *Report-Oblivious Spot-checking* (ROS) mechanism spot checks every student with fixed probability x , regardless of the students' reports.

Definition 3.4 (RSS Mechanism). A *Report-Sensitive Spot-checking* (RSS) mechanism spot checks every student with probability that can depend on all the students' reports.

Definition 3.5 (RSUS Mechanism). A *Report-Sensitive, Uniform Spot-checking* (RSUS) mechanism ensures that whenever one student is spot checked, all are spot checked.

3.2 Assumptions

Our first assumption is that a student, upon observing a signal, expects any other grader (the TA or another student) to be strictly more likely to observe the same signal than the opposite. This assumption is needed to ensure that students are strictly incentivized to report honestly in ROS mechanisms. Our second assumption is that a student being spot checked with probability 1 prefers to be truthful than to report an arbitrary signal without effort. In effect, this assumption can be understood as saying that rewards are large enough to outweigh the cost of effort.

4 SUMMARY OF THE RESULTS

We focus on minimizing the need for *TA workload*. For a DSIC spot-checking mechanism, the TA workload is the probability with which the TA needs to provide a signal, assuming all students are truthful.

Our first technical result is a characterization of the optimal DSIC RSS mechanism. We define a class of mechanisms (PRSS) and then show that the minimum TA workload is always achieved by a DSIC PRSS mechanism.

Definition 4.1 (PRSS Mechanism). A *Personal-Report-Sensitive Spot-checking mechanism*, or *PRSS mechanism*, spot checks each student with a probability that only depends on the student's own report.

An immediate consequence of this characterization is that, as long as the two signals appear with different ex ante probabilities, the optimal DSIC RSS mechanism spot checks strictly less than the any DSIC ROS mechanism.

The proof consists of two steps. In Step 1, we reason about a convex optimization problem that minimizes the TA workload but relaxes all of the DSIC constraints except those that incentivize the truthful strategy when all other students make no effort. We identify an optimal solution to this problem making use of the convexity of the objective function. Then, in Step 2, we show that this solution in fact gives rise to a DSIC PRSS mechanism. Since it is an optimal solution with most DSIC constraints relaxed, it is also optimal when one enforces all constraints.

Our second result shows limitation of RSUS mechanisms. RSUS mechanisms have the intuitively appealing property that the TA's signal is never "wasted": whenever TA is asked to provide a signal, it is used to spot check all students. We show that such mechanisms are nevertheless outperformed in general by our optimal mechanism, which is PRSS. We prove an even stronger result, showing that the optimal DSIC PRSS mechanism outperforms RSUS mechanisms even when the latter is subject to a weaker solution concept, which requires the truthful strategy to be utility maximizing for a student *as long as any other student who examines the assignment always reports the observed signal*.

REFERENCES

- [1] 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.
- [2] Luca De Alfaro and Michael Shavlovsky. 2014. CrowdGrader: A tool for crowdsourcing the evaluation of homework assignments. In *Proceedings of the 45th ACM Technical Symposium on Computer Science Education*. ACM, 415–420.
- [3] Luca De Alfaro, Michael Shavlovsky, and Vassilis Polychronopoulos. 2016. Incentives for truthful peer grading. *arXiv preprint arXiv:1604.03178* (2016).
- [4] Boi Faltings, Jason J Li, and Radu Jurca. 2012. Eliciting truthful measurements from a community of sensors. In *3rd IEEE International Conference on the Internet of Things*. IEEE, 47–54.
- [5] Alice Gao, Andrew Mao, Yiling Chen, and Ryan P. Adams. 2014. Trick or treat: putting peer prediction to the test. In *Proceedings of the 15th ACM Conference on Economics and Computation*. ACM, 507–524.
- [6] Alice Gao, James R Wright, and Kevin Leyton-Brown. 2016. Incentivizing evaluation via limited access to ground truth: Peer-prediction makes things worse. *arXiv preprint arXiv:1606.07042* (2016).
- [7] Radu Jurca and Boi Faltings. 2005. Enforcing truthful strategies in incentive compatible reputation mechanisms. In *International Workshop on Internet and Network Economics*. Springer, 268–277.
- [8] Radu Jurca and Boi Faltings. 2009. Mechanisms for making crowds truthful. *Journal of Artificial Intelligence Research* 34 (2009), 209–253.
- [9] Vijay Kamble, David Marn, Nihar Shah, Abhay Parekh, and Kannan Ramachandran. 2015. Truth Serums for Massively Crowdsourced Evaluation Tasks. *arXiv preprint arXiv:1507.07045* (2015).
- [10] Yuqing Kong, Katrina Ligett, and Grant Schoenebeck. 2016. Putting Peer Prediction Under the Micro (economic) scope and Making Truth-telling Focal. In *International Conference on Web and Internet Economics*. Springer, 251–264.
- [11] Michael R Merrifield and Donald G Saari. 2009. Telescope time without tears: A distributed approach to peer review. *Astronomy & Geophysics* 50, 4 (2009), 4–16.
- [12] Nolan Miller, Paul Resnick, and Richard Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* 51, 9 (2005), 1359–1373.
- [13] Dražen Prelec. 2004. A Bayesian truth serum for subjective data. *Science* 306, 5695 (2004), 462–466.
- [14] Goran Radanovic and Boi Faltings. 2013. A robust Bayesian truth serum for non-binary signals. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence*. 833–839.
- [15] Goran Radanovic and Boi Faltings. 2014. Incentives for truthful information elicitation of continuous signals. In *Proceedings of the 28th AAAI Conference on Artificial Intelligence*. 770–776.
- [16] Blake Riley. 2014. Minimum truth serums with optional predictions. In *Proceedings of the 4th Workshop on Social Computing and User Generated Content*.
- [17] Victor Shnayder, Arpit Agarwal, Rafael Frongillo, and David C Parkes. 2016. Informed truthfulness in multi-task peer prediction. In *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 179–196.
- [18] Luis Von Ahn and Laura Dabbish. 2004. Labeling images with a computer game. In *Proceedings of the SIGCHI conference on Human factors in computing systems*. ACM, 319–326.
- [19] Luis Von Ahn and Laura Dabbish. 2008. Designing games with a purpose. *Commun. ACM* 51, 8 (2008), 58–67.
- [20] Bo Waggoner and Yiling Chen. 2014. Output agreement mechanisms and common knowledge. In *2nd AAAI Conference on Human Computation and Crowdsourcing*.
- [21] Wanyuan Wang, Bo An, and Yichuan Jiang. 2018. Optimal Spot-Checking for Improving Evaluation Accuracy of Peer Grading Systems. In *32nd AAAI Conference on Artificial Intelligence*. 833–840.
- [22] Jens Witkowski, Yoram Bachrach, Peter Key, and David C Parkes. 2013. Dwelling on the negative: Incentivizing effort in peer prediction. In *1st AAAI Conference on Human Computation and Crowdsourcing*.
- [23] Jens Witkowski and David C Parkes. 2012. A robust Bayesian truth serum for small populations. In *26th AAAI Conference on Artificial Intelligence*. 1492–1498.
- [24] James R Wright, Chris Thornton, and Kevin Leyton-Brown. 2015. Mechanical TA: Partially automated high-stakes peer grading. In *Proceedings of the 46th ACM Technical Symposium on Computer Science Education*. ACM, 96–101.