Eliciting Expertise without Verification

YUQING KONG, University of Michigan, USA

GRANT SCHOENEBECK, University of Michigan, USA

A central question¹ of crowdsourcing is how to elicit expertise from agents. This is even more difficult when answers cannot be directly verified. A key challenge is that sophisticated agents may strategically withhold effort or information when they believe their payoff will be based upon comparison with other agents whose reports will likely omit this information due to lack of effort or expertise.

Our work defines a natural model for this setting based on the assumption that *more sophisticated agents know the beliefs of less sophisticated agents*.

We then provide a mechanism design framework for this setting. From this framework, we design several novel mechanisms, for both the single and multiple tasks settings, that (1) encourage agents to invest effort and provide their information honestly; (2) output a correct “hierarchy” of the information when agents are rational.

CCS Concepts: • **Theory of computation** → **Algorithmic mechanism design**; • **Information systems** → **Incentive schemes**;

Additional Key Words and Phrases: Peer prediction, information theory

1 INTRODUCTION

Crowdsourcing, outsourcing tasks to a crowd of workers (e.g. Amazon Mechanical Turk, peer grading for massive open online courses, scholarly peer review, and Yahoo answers), is a fast, cheap, and effective method for performing simple tasks even at large scales. To attract a large number of workers, crowdsourcing is usually open to the public rather than just professional experts. This makes crowdsourcing a challenge for more complicated tasks where agents with different levels of expertise may honestly report different answers. For example, in peer grading, some students may be incapable of accurate grading, missing errors other students point out.

One solution is to spot-check random tasks or equivalently to insert “gold-standard” tasks to which the mechanism already knows the answer. However, such techniques are expensive and even impossible for tasks where the agents’ opinions can be very subjective (e.g. peer grading essay assignments) or individualized (e.g. experience at a restaurant).

¹ A full version of the paper is available at <https://arxiv.org/abs/1802.08312>.

This work is supported by the National Science Foundation, under grant CAREER#1452915, CCF#1618187 and AitF#1535912. Authors’ addresses: Yuqing Kong, University of Michigan, Ann Arbor, USA, yuqkong@umich.edu; Grant Schoenebeck, University of Michigan, Ann Arbor, USA, schoeneb@umich.edu.

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

© 2018 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM EC'18, June 18–22, 2018, Ithaca, NY, USA. ACM ISBN 978-1-4503-5829-3/18/06...\$15.00

<https://doi.org/10.1145/3219166.3219172>

1.1 Our setting

This work is about the “peer-prediction” setting where such verification is unavailable or undesirable. Already, a large line of work in the peer prediction (PP) area [3, 6, 8, 12–24, 27, 28, 30–34, 36, 39–42] designed mechanisms for agents with similar expertise (without verification).

These papers study various settings. In the *single-task* setting [27, 28] agents are asked one multiple-choice question but typically share a common prior over the responses; in the *multi-tasks* setting [6, 20, 36] agents are assigned a batch of a priori similar multiple-choice questions. We consider both of these settings.

In the *detail-free* setting, the mechanism is not required to know information about the (common) prior distribution. Our mechanisms are generally in this setting, but some of them do require information about the different levels of expertise/effort agents can exhibit.

However, the aforementioned peer-prediction papers do not consider settings where

- (1) **Agents have different levels of expertise** or
- (2) **A lack of effort can systemically bias agents’ reports.**

The following two tasks exemplify settings 1) and 2) respectively:

Example 1.1. Which state (from a list of all 50 states) in the United States of America is closest to Africa? (Single-task)

Example 1.2. Peer grading several essays by providing a grade from the set {1, 2, 3, 4, 5}. (Multi-tasks)

In the first example, an agent can guess randomly (no effort), look up the correct answer (full effort), or guess at the correct answer (partial effort). Most people will guess Florida, even though experts will know the correct answer is Maine. Thus differing levels of expertise yield different answers.

In the second example, a student can, instead of carefully grading (full effort) or assigning a random grade (zero effort), quickly check the name of the top of the paper and spot check the grammar (partial effort). Thus partial effort can systematically bias agents: consider an essay from a top student in an impeccable pose, but which contains large conceptual errors. Here partial effort can give some information about the correct answer, but also enable agents to “coordinate” on an incorrect answer.

Gao, Wright, and Leyton-Brown [9] show that the effects of the settings 1) and 2) are devastating to previous peer-prediction mechanisms, which generally fail in motivating the agents to invest effort for “expensive signals” when “cheap signals” (that ensure agreement and may even be correlated with the sought signal) exist. The main (very high-level) idea behind previous peer-prediction mechanisms can be understood as a “clever majority vote”, every agent is paid according to a specific similarity between her and her peer. Thus, they point out that in the peer-grading example, coordinating on just checking the grammar can guarantee good agreement with other agents, but with substantially reduced effort.

In fact, Gao et al point out that things are likely even worse than this. If the cheap signals correlate more than the expensive signals, then the peer-prediction techniques incentivize agents to not report the true answer, but instead focus on cheap signals! For example, in the essay grading above, it is likely that assessments of grammatical correctness will agree more than assessments of overall essay quality. Because of this, peer-prediction mechanisms will pay agents more overall for lower-quality information. In Example 1.1, even if agents know the answer is Maine, they may report Florida, expecting that most others will do likewise.

Such behavior undermines the goal of applying crowdsourcing to increasingly complex tasks, and, in fact, undercuts any application of crowdsourcing to perform any task where the answers

are not “common knowledge.” The field must overcome this key challenge of rewarding rather than suppressing expertise in order to begin the project of expanding crowdsourcing beyond simple labeling tasks.

1.2 Key Insight and Assumption

Previous peer-prediction mechanisms treat all information (cheap/expensive) equally, and thus agents lack an incentive to invest the effort to obtain expensive signals. Moreover, even when an expert can easily obtain the expensive signal, the previous mechanism discourage her from providing it when she believes the non-experts will disagree.

A successful mechanism must break the symmetry between weak and expensive signals and between expert and non-expert signals. We propose the following natural assumption which will allow a mechanism to break this symmetry.

ASSUMPTION 1.3. *Agents with high effort or expertise know the beliefs of agents with less effort or less expertise.*

We can see that this assumption is very natural in Example 1.1 and Example 1.2. Agents who look up or know the answer in Example 1.1 also know most people will answer “Florida.” Agents that carefully grade an essay can also approximate the score of an agent who spends very little effort. We will define a hierarchical information structure (Section 3) to naturally capture Assumption 1.3.

Our mechanisms solicit not only agents’ own opinions but also their predictions for the opinions of the other agents who have less information. This differs with the previous peer prediction mechanisms which ask agents to provide their predictions for *all* other agents’ opinions. For example, in the peer-grading example, we might ask agents to report their own evaluation, and to *optionally* report one or more low-effort / low expertise evaluations (e.g. scores based on grammar, thesis statement, student name, naive reading, etc).

1.3 Our contribution

Modeling: We define a hierarchical information structure (the higher level information is more valuable and harder to obtain) model that captures the hierarchical relationship between the different expertise;

Mechanism design: We employ information theory tools to provide a mechanism design framework for the aforementioned hierarchical setting—Hierarchical Mutual Information Paradigm (HMIP)—which 1) pays the higher level information more and encourages agents to invest effort for different levels of information according to their own abilities 2) incentivizes agents to report honestly; and 3) only pays agents for information they actually learn. We apply our mechanism design framework to create the following mechanisms:

Multi-HMIM which works in the multi-task setting even for a small number of tasks but requires the mechanism to know the hierarchical information structure.

Learning-based Multi-HMIM which works in the multi-task setting even when the mechanism does not know the hierarchical information structure; however requires a large number of tasks.

Single-HMIM which works in the single-task setting.

Algorithmic: In the multiple task setting where the mechanism does not know the hierarchical information structure, we design an algorithm which, given agent reports as input, outputs the hierarchical structure of the information as well as the most valuable information. Note that our algorithm is robust against agents with limited (or no) information reporting biased (or totally random and useless) information. This is used as part of the *Learning-based Multi-HMIM* mechanism.

1.4 Road Map

The remainder of this section discusses related work; Section 2 lays out the mathematical tools we use and can be referred back to when necessary. Section 3 introduces both our model, the hierarchical information structure, and our mechanism design framework, the hierarchical mutual information paradigm (HMIP). Section 4 and section 5 apply the HMIP mechanism design framework into multi-task setting and single-task setting respectively.

In the multi-task setting (Section 4), we consider two situations: 1) the mechanism knows the information structure shared by agents but each agent is assigned a small number of tasks; 2) the mechanism does not know the information structure but each agent is assigned a large number of tasks such that the mechanism can learn the information structure. In the single-task setting (Section 5), we assume the mechanism know the information structure.

Section 6 concludes by discussing some potential applications and the robustness of our mechanisms.

1.5 Related Work

Model perspective. Prior work has modeled heterogeneous expertise where different agents receive a different number of signals [12] or expertise is embedding in several dimensions [7, 26, 37, 43]; however in these works lower expertise/effort along with a certain dimension only leads to a more noisy signal. In contrast, our model allows such signals to be systematically biased.

Mechanism design perspective. The most related work with the current paper is Prelec et al. [29] which uses Bayesian Truth Serum [28] to incentivize agents to report their signal and selects the most surprising signal (measured by occurring more than its average prediction) as the final answer. McCoy and Prelec [26] follow Prelec et al. [29] to propose a probabilistic model to learn the expertise of agents. Riley [35] compares the peer prediction decision rule (similar to Prelec et al. [29]) and the majority vote rule and exhibits cases where each outperforms the other. The current paper differs with Prelec et al. [29] in the model and assumptions as well as the possible applications. Prelec et al. [29] only focus on the single-task setting and assume that agents receive the signals endogenously (without effort). In contrast, this paper considers both single and multiple task settings and the model used in this paper handles both exogenous and indigenous signals.

The mechanism design framework in the current paper extends the information theoretic framework proposed in Kong and Schoenebeck [20]. Agarwal et al. [1] propose a mechanism that works for the heterogeneous participants in the multi-task setting. Mandal et al. [25] consider the setting with heterogeneous tasks. They all do not assume the hierarchy of the information and cannot be applied to identify and elicit expertise.

Algorithmic perspective. Several works [7, 10, 37, 43] provide clever methods to learn the expertise as well as the ground truth of the crowdsourcing tasks. The algorithm in the current paper differs in two main aspects: (1) The current paper uses a different expertise model which can successfully capture the possibly hierarchical relationship between different information/expertise as well as the most valuable information; (2) the current paper combines the algorithm with an incentive mechanism that endogenously controls the quality and structure of the input, rather than making exogenous assumptions about the quality of the input.

2 MECHANISM DESIGN TOOLS

We use two key information theory ingredients in designing information elicitation mechanisms. The first ingredient is f -mutual information $MI^f(X; Y)$ which measures the amount of information crossing two random variables X, Y . For example, if X is independent with Y —no information

crosses X and Y , $MI^f(X; Y) = 0$. The second ingredient is proper scoring rule $PS(x, \mathbf{p})$ which measures the accuracy of the prediction $\mathbf{p}$ even we only have one sample x of the outcome X .

Both two ingredients have the information monotonicity property. If the information is measured by f -mutual information, any “data processing” on either of the random variables will decrease the amount of information crossing them. If the accuracy of a forecast is measured by a proper scoring rule, more information implies a more accurate forecast.

2.1 f -mutual information

f -divergence [2, 5] $D_f : \Delta_\Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$ is a non-symmetric measure of the difference between distribution $\mathbf{p} \in \Delta_\Sigma$ and distribution $\mathbf{q} \in \Delta_\Sigma$ and is defined to be $D_f(\mathbf{p}, \mathbf{q}) := \sum_{\sigma \in \Sigma} \mathbf{p}(\sigma) f\left(\frac{\mathbf{q}(\sigma)}{\mathbf{p}(\sigma)}\right)$ where $f(\cdot)$ is a convex function and $f(1) = 0$. Two commonly used f -divergences are KL divergence and total variation distance. Now we start to introduce f -mutual information.

Given two random variables X, Y , let $\mathbf{U}_{X,Y}$ and $\mathbf{V}_{X,Y}$ be two probability measures where $\mathbf{U}_{X,Y}$ is the joint distribution of (X, Y) and $\mathbf{V}$ is the product of the marginal distributions of X and Y . Formally, for every pair of (x, y) ,

$$\mathbf{U}_{X,Y}(X = x, Y = y) = \Pr[X = x, Y = y] \quad \mathbf{V}_{X,Y}(X = x, Y = y) = \Pr[X = x] \Pr[Y = y].$$

If $\mathbf{U}_{X,Y}$ is very different with $\mathbf{V}_{X,Y}$, the mutual information between X and Y should be high since knowing X changes the belief for Y a lot. If $\mathbf{U}_{X,Y}$ equals to $\mathbf{V}_{X,Y}$, the mutual information between X and Y should be zero since X is independent with Y . Intuitively, the “distance” between $\mathbf{U}_{X,Y}$ and $\mathbf{V}_{X,Y}$ represents the mutual information between them.

Definition 2.1 (f -mutual information [20]). The f -mutual information between X and Y is defined as

$$MI^f(X; Y) = D_f(\mathbf{U}_{X,Y}, \mathbf{V}_{X,Y})$$

where D_f is f -divergence.

Definition 2.2 (Conditional f -mutual information [20]). Given three random variables X, Y, Z , we define $MI^f(X; Y|Z)$ as

$$\sum_z \Pr[Z = z] MI^f(X; Y|Z = z)$$

where $MI^f(X; Y|Z = z) := MI^f(X'; Y')$ where $\Pr[X' = x, Y' = y] = \Pr[X = x, Y = y|Z = z]$.

Two examples of f -mutual information are Shannon mutual information [4] (Choosing f -divergence as KL divergence) and $MI^{tvd}(X; Y) := \sum_{x,y} |\Pr[X = x, Y = y] - \Pr[X = x] \Pr[Y = y]|$ (Choosing f -divergence as Total Variation Distance).

LEMMA 2.3 (GENERAL DATA PROCESSING INEQUALITY / INFORMATION MONOTONICITY [20]). When f is convex, f -mutual information MI^f is symmetric $MI^f(X; Y) = MI^f(Y; X)$; non-negative and satisfies data processing equality: for any transition probability $M \in \mathbb{R}^{|\Sigma_X| \times |\Sigma_X|}$, when Y is independent with $M(X)$ conditioning on X , $MI^f(M(X); Y) \leq MI^f(X; Y)$.

2.2 Proper scoring rules

Informally, a scoring rule measures the accuracy of the forecasts. Formally, a scoring rule [11, 38] $PS : \Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$ takes in a signal $x \in \Sigma$ and a distribution over signals $\delta_\Sigma \in \Delta_\Sigma$ and outputs a real number. A scoring rule is *proper* if, whenever the first input is drawn from a distribution δ_Σ , then δ_Σ will maximize the expectation of PS over all possible inputs in Δ_Σ to the second coordinate. A scoring rule is called *strictly proper* if this maximum is unique. We will assume throughout that the scoring rules we use are strictly proper. Slightly abusing notation, we can extend a scoring rule to be $PS : \Delta_\Sigma \times \Delta_\Sigma \rightarrow \mathbb{R}$ by simply taking $PS(\delta_\Sigma, \delta'_\Sigma) = \mathbb{E}_{x \leftarrow \delta_\Sigma}(x, \delta'_\Sigma)$.

3 MODEL AND MECHANISM DESIGN FRAMEWORK

In this section, we will define the hierarchical information structure and provide a mechanism design framework that helps design mechanisms which elicit the hierarchical information. Section 3.1 defines the information model; Section 3.2 defines our mechanism framework; and Section 3.3 analyzes the framework. We will use the peer grading process (Figure 1) as a running example to throughout this section.

Fig. 1. An illustration of the hierarchical information structure in the peer grading process.

3.1 Hierarchical Information Structure

There are n agents and one task. The agents have a finite set M of methods to perform on the task based on the task's attributes $\mathbf{a} \in A$ where $\mathbf{a}$ is a random (possibly high dimensional) vector drawn from a distribution $Q_A \in \Delta_A$. Each method $m : A \mapsto \Sigma_m$ maps the attributes $\mathbf{a} \in A$ to a signal $m(\mathbf{a})$ from a finite set Σ_m . We now introduce our peer grading example.

10 evaluators are asked to judge one essay. The essay has eight possible attributes: $\mathbf{a} = (q_i, w_j, l_k), i, j, k \in \{0, 1\}$. (q_1, w_0, l_1) means the essay has (good quality, bad writing, long length); (q_0, w_1, l_0) means the essay has (bad quality, good writing, short length). The distribution over the attributes space Q_A is defined as:

$Q_A((q_0, w_0, *))$	$Q_A((q_0, w_1, *))$	$Q_A((q_1, w_0, *))$	$Q_A((q_1, w_1, *))$
0.4	0.1	0.1	0.4

With this distribution, an essay of good quality usually has good writing as well. Moreover, we assume the essay's length is independent with the essay's quality and writing and an essay has long length with probability 0.5. That is: $Q_A((q_i, w_j, l_1)) = Q_A((q_i, w_j, *)) * 0.5$, $Q_A((q_i, w_j, l_0)) = Q_A((q_i, w_j, *)) * 0.5$.

Each evaluator can perform three methods: $m_l(\mathbf{a})$, $m_w(\mathbf{a})$, and $m_q(\mathbf{a})$ which are, respectively, (possibly noisy) signals about the essay's length; writing style and grammar; and quality. $\Sigma_l = \Sigma_w = \Sigma_q = \{\odot, \ominus\}$.

We define $\psi_i^m(\mathbf{a})$ as agent i 's received output by performing m on attributes $\mathbf{a}$. Different agents may receive different signals by performing the same method on the same attributes. But we assume the distribution is symmetric/homogeneous in the sense that for any permutation $\pi : [n] \mapsto [n]$, the probability that $\psi_1^m(\mathbf{a}) = \sigma_1, \psi_2^m(\mathbf{a}) = \sigma_2, \dots, \psi_n^m(\mathbf{a}) = \sigma_n$ equals the probability that $\psi_{\pi(1)}^m(\mathbf{a}) = \sigma_1,$

$\psi_{\pi(2)}^m(\mathbf{a}) = \sigma_2, \dots, \psi_{\pi(n)}^m(\mathbf{a}) = \sigma_n$. We also assume that each agent performs methods independently (see (1)). When the attributes $\mathbf{a}$ is drawn from a distribution Q_A , we can define Ψ_i^m as agent i 's received output by performing m on a random attributes $\mathbf{a}$ that is drawn from a distribution Q_A . Analogously, we define a random variable Ψ_{-i}^m as an arbitrary agent $j \neq i$'s received output by performing m on a random attributes $\mathbf{a}$ that is drawn from a distribution Q_A . This definition is well-defined since we have assumed the distribution is symmetric. We define prior Q as a joint distribution over all $\{\Psi_i^m\}_{i \in [n], m \in M}$.

Conditioning on the attributes of the essay $\mathbf{a} = (q_i, w_j, l_k), i, j, k \in \{0, 1\}$, for each method m , each agent will receive $\psi_i^m(\mathbf{a}) = \odot$ with probability $p_{m,\mathbf{a}}$ independently by performing m . That is, agents' received signals by performing m is a Binomial distribution $B(n = 10, p_{m,\mathbf{a}})$.^a

	good quality essay ^b	bad quality essay
$\Pr[m_q(\mathbf{a}) = \odot]$	70%	30%

This means conditioning on the essay having good quality, the distribution over agents' received quality signals by performing m_q is $Q_{m_q, (q_1, *, *)} = B(10, 0.7)$; while conditioning on the essay having bad quality, the distribution is $Q_{m_q, (q_0, *, *)} = B(10, 0.3)$. Similarly, we have

	good writing essay	bad writing essay
$\Pr[m_w(\mathbf{a}) = \odot]$	90%	10%
	long essay	short essay
$\Pr[m_l(\mathbf{a}) = \odot]$	100%	0%

Note that the cheap length signal is noiseless. We also assume that fixing the attributes, every agent performs the different methods independently. That is, when $\mathbf{a} = (q_1, w_1, l_1)$

$$\Pr\left(\Psi_i^{m_l}(\mathbf{a}) = \odot, \Psi_i^{m_w}(\mathbf{a}) = \odot, \Psi_i^{m_q}(\mathbf{a}) = \odot\right) = 0.7 * 0.9 * 1. \quad (1)$$

With the above set up, the probability that agent i receives a $\odot$ writing signal and agent j receives a $\odot$ quality signal will be

$$\Pr[\Psi_i^{m_w} = \odot, \Psi_j^{m_q} = \odot] = 0.4 * 0.1 * 0.3 + 0.1 * 0.3 * 0.9 + 0.1 * 0.7 * 0.1 + 0.4 * 0.7 * 0.9 = 0.298.$$

^aTo give a concrete example, we use the Binomial distribution here. In fact, we only need the distribution to be symmetric.

^bThis means $\mathbf{a} = (q_1, *, *)$

We define a partial order on the methods. We say $m_1 \geq m_2$ —the level of m_1 is higher than that of m_2 —if method m_1 cannot be performed without performing m_2 . By $m_1 > m_2$ we mean $m_1 \geq m_2$ but $m_2 \not\geq m_1$. Note that the partial order $>$ is transitive— $m_1 > m_2, m_2 > m_3 \Rightarrow m_1 > m_3$. Each agent i needs effort $h_i(m) > 0$ to perform method m and when she spends effort $h_i(m)$ to perform m , the methods that have lower levels than m are performed as well without additional effort. We assume, as is natural, that $h_i(m)$ is an increasing function, that is, $h_i(m_1) \geq h_i(m_2)$ when $m_1 \geq m_2$. The higher the level of the method an agent performs, the more effort she must invest. However, it may be the case that some agents (low cost agents) can perform methods more economically than others (high cost agents). The partial order definition is essentially our key assumption (Assumption 1.3).

$m_q > m_w > m_a$. Among the 10 evaluators, there are 2 low cost evaluators who need 1, 2, 5 effort to perform m_l, m_w, m_q respectively. There are 8 high cost evaluators who need 1, 4, 10 effort to perform m_l, m_w, m_q respectively (Figure 1). Based on the partial order definition, when

an evaluator spends sufficient effort to perform m_q and obtains the quality signal, she also obtains the length and writing signals without additional effort, which is natural in real life.

We assume agents share a hierarchical information structure and allow agents to have different priors Q^2 . We will design mechanisms that incentivize agents to invest efforts based on their costs and report honestly.

3.2 Mechanism Design Framework

We start by introducing the formal definition of a mechanism.

Definition 3.1 (Mechanism). We define a *mechanism* $\mathcal{M}$ for n agents as a tuple $\mathcal{M} := (R, S)$ where R is a set of all possible reports the mechanism allows, and $S : R^n \mapsto \mathbb{R}^n$ is a mapping from all agents' reports to each agent's payment.

We will extend the Mutual Information Paradigm [20] to our Hierarchical Mutual Information Paradigm that handles the hierarchical information structure.

Mutual Information Paradigm [20]. Kong and Schoenebeck [20] provide a Mutual Information Paradigm for Peer Prediction Mechanisms: each agent i is paid the mutual information between her information and her peers' information—

$$MI(\text{her information; her peers' information}).$$

By picking f -mutual information MI^f that satisfies data processing equality (Section 2.1), no agent can obtain strict benefit by lying since intuitively the amount of information each agent has will not increase no matter what kind of strategy she applies to her information.

We can naturally extend the Mutual Information Paradigm for Peer Prediction Mechanisms to the hierarchical model by paying agent i

$$MI^f(\text{her information; } \{\Psi_{-i}^m\}_{m \in M})^3.$$

This idea has a severe drawback: sometimes low level information has very large correlation with the high level information. In this case, $MI^f(\text{her information; } \{\Psi_{-i}^m\}_{m \in M})$ will pay low level information nearly as much as high level information; and so agents will lack incentive to perform high level methods.

To solve the above problem, we pay agents method by method. For each m , we only value the “information gain” in the sense that we pay each agent i the mutual information between her information and the method m 's information *conditioning on* the information output by the methods are lower than m .

Formally, we chose a payment scale α_m for each m and pay each agent i

$$\sum_m \alpha_m MI^f(\text{her information; } \Psi_{-i}^m | \{\Psi_{-i}^{m'}\}_{m' < m}).$$

In our actual paradigm, we hope to pay each agent i using the above payment when the mechanism has access to all levels of honest information provided by other agents.

In the peer grading example, the information about the writing style / grammar may already have a very high correlation with the quality of the essay. With the above concrete set up,

²To ease the presentation of the example, in our peer grading example, we assume agents share the same prior Q .

³Recall that $(\Psi_1^m, \Psi_2^m, \dots, \Psi_n^m)$ are the random signals agents receive by performing method m on the same random attributes.

we are ready to calculate the (conditional) Shannon mutual information (Euler number base) between agent i 's received signals and agent j 's received signals. For example, the 2×2 entry is the mutual information between agent i 's received length signal, writing signal by performing method m_w and agent j 's writing signal, conditioning on agent j 's length signal, which is 0.2259. We calculate the values by first calculating the joint distribution over 6 random variables—agent i 's length, writing, quality signals and agent j 's length, writing, quality signals. We show the values in the following table and defer the calculation to the full version. According to the information monotonicity, for each column, the values increase from bottom to top.

$MI(\cdot; \cdot)$ agent i 's	agent j 's			
	length	writing length ^a	quality writing, length	length, writing, quality ^b
length, writing, quality	0.6931	0.2259	0.0115	0.9305
length, writing	0.6931	0.2218	0.0041	0.9190
length	0.6931	0	0	0.6931

Even though performing the quality method provides the information that has the highest mutual information 0.9305 with other agents' information, performing writing method already outputs information that has $0.9190 \approx 0.9305 * 0.98$ mutual information with other agents' information.

In this case, what we really value is the additional quality of information after conditioning on the information of cheap signals like writing style / grammar. In other words, we value the information about an essay which has a high quality but is written carelessly (or low quality but impeccable prose).

Each agent, performing the writing method only has 0.0041 mutual information with other agents' quality signal conditioning on other agents' writing and length signals while performing the quality method has $0.0115 \approx 0.0041 * 2.80$ conditional mutual information.

Looking ahead, we seek to pay each evaluator i by:

$$\begin{aligned} & \alpha_l MI^f(\text{her information; agent } j\text{'s length signal}) \\ & + \alpha_w MI^f(\text{her information; agent } j\text{'s writing signal} | \text{agent } j\text{'s length signal}) \\ & + \alpha_q MI^f(\text{her information; agent } j\text{'s quality signal} | \text{agent } j\text{'s length and writing signal}). \end{aligned} \quad (2)$$

where α_q is set to be rather larger than α_l and α_w .

^a $x|y$ means x conditioning on y .

^bSince we use Shannon mutual information which satisfies chain rule, the last column is the sum of the previous columns.

Hierarchical Mutual Information Paradigm (HMIP($MI^f, \{\alpha_m\}_m$)). We now present our hierarchical Mutual Information Paradigm. We emphasize that this is not a mechanism that can be run. Instead we engage in the wishful thinking that the reports of the agents are distributions rather than draws from the distribution. Of course, this will never happen. Nonetheless, we will show that using the HMIP paradigm we can design actual mechanisms in both the multiple-task setting (Section 4) and the single-task setting (Section 5).

The paradigm requires as parameters a payment scale $\alpha_m \in \mathbb{R}_{\geq 0}$ for each method m .

Report For each agent i , for each $m \in M$, she is asked to optionally provide the random signal Ψ_i^m . We denote the set of methods whose outputs are reported by agent i as M_i and the actual random signal she reports for each $\ell \in M_i$ as $\hat{\Psi}_i^\ell$.

Payment/Information Score We define M_{-i} as $\bigcup_{j \neq i} M_j$. For each $m \in M_{-i}$, we arbitrarily pick an agent $j \neq i$ who provides method m 's output and denote his report for method m 's output as $\hat{\Psi}_{-i}^m$.

Agent i is paid by her information score

$$\sum_{m \in M_{-i}} \alpha_m MI^f(\{\hat{\Psi}_i^\ell\}_{\ell \in M_i}; \hat{\Psi}_{-i}^m | \{\hat{\Psi}_{-i}^{m'}\}_{m' < m, m' \in M_{-i}})$$

Resolving “wishful thinking”. HMIP pays agents according to the information measure. The calculation of the information measure requires the knowledge of the prior, i.e., the joint distribution which is unrealistic in practice. To resolve this “wishful thinking”, a key observation is that paying agents an unbiased estimator of the information measure is sufficient when we assume agents are expected utility maximizers. To construct an unbiased estimator of the information measure using agents' reports, different settings have different techniques. In the multi-task setting, either we ask a large number of questions to estimate the prior and use the prior to calculate the information measure (Learning-based multi-HMIM), or we ask a small number questions but require the knowledge of the structure and use a special f -mutual information, MI^{tvd} (Multi-HMIM). In the single-task setting (Single-HMIM), we ask agents their posteriors (e.g. what percentage of your peers say yes?) and construct the estimator using both the first order information (e.g., Y/N) and the second order information (e.g., 80% Yes). Thus, although all proposed mechanisms are based on the HMIP, they are all detail free in the sense that they do not need any priori knowledge of the distributions (nor wishful thinking).

3.3 Analysis of HMIP

For each agent i , we define her utility as her payment minus her effort.

Definition 3.2 (Strategy). We define the *effort strategy* of each agent i as a mapping e_i from her priors to a probability distribution over the methods she will perform. We define the *report strategy* of each agent i as a mapping s_i from her received information to a probability distribution over R .

Definition 3.3 (Amount of information in HMIP). In HMIP, for agent i , the amount of information acquired with method m_i is defined as

$$AOI(m_i, \text{HMIP}(MI^f, \{\alpha_m\}_m)) := \sum_{m \in M} \alpha_m MI^f(\{\Psi_i^\ell\}_{\ell \leq m_i}; \Psi_{-i}^m | \{\Psi_{-i}^{m'}\}_{m' < m}).$$

We have already give the example of the amount of information in (2). Later in the proof of Theorem 3.7, we will see the amount of information is also the optimal payment of agent i who performs method m_i when HMIP has access to all levels of honest signals reported by other agents.

An especially desirable strategy in HMIP is a *prudent strategy*. Informally, agents play a prudent strategy if they (a) choose the method they perform to maximize their utility—trading off the amount of information acquired with the effort it costs; (b) report all received information honestly.

Definition 3.4 (Prudent strategy in HMIP). For each agent i , we say she plays a prudent strategy in $\text{HMIP}(MI^f, \{\alpha_m\}_m)$ if she chooses to (a) perform method m_i^* such that

$$m_i^* \in \arg \max_{m_i} \left(AOI(m_i, \text{HMIP}(MI^f, \{\alpha_m\}_m)) - h_i(m_i) \right); \text{ and}$$

(b) reports all received information honestly.

Definition 3.5 (Truthful strategy in HMIP). We say an agent plays truthful strategy if she always reports her received information honestly.

A truthful strategy is a special report strategy. An agent can play any effort strategy and truthful strategy simultaneously. If an agent invests no effort and reports nothing or meaningless information, she is still considered as playing truthful strategy.

Mechanism design goals. A mechanism $\mathcal{M}$ is (strictly) *potent* if for each agent, when she believes everyone else plays their prudent strategy, she can (strictly) maximize her expected utility by playing a prudent strategy as well. A mechanism $\mathcal{M}$ is *dominant truthful* if for each agent, regardless of other agents' strategies, she can maximize her expected utility by playing a pure effort strategy and truthful strategy.

The dominant truthful property is incomparable with the potent property. A flat payment scheme is dominant truthful but not potent since investing no effort and reporting nothing is also considered as a pure effort and truthful strategy. The potent property is desirable since it encourages low cost agents to invest high level effort and high cost agents to invest low level effort, and incentivizes them to report honestly as well.

In order to design potent mechanism, the coefficients $\{\alpha_m\}_m$ should be chosen appropriately. If the coefficients $\{\alpha_m\}_m$ are too low, no agent will be incentivized to invest the highest level method. In this case, the mechanism cannot access all levels of information to encourage agents to play prudent strategy. Thus, we will assume the coefficients are chosen such that it is worthwhile for at least two agents to invest the highest level method.

We say a method m is *maximal* if there does not exist $m' \neq m \in M$ such that $m' > m$.

Definition 3.6 (potent coefficients for HMIP). Given the priors $\{Q_m\}_m$, we say the coefficients $\{\alpha_m\}_m$ are potent for $\text{HMIP}(M^f, \{\alpha_m\}_m)$ if given the coefficients $\{\alpha_m\}_m$, for every maximal m , there exists at least **two** agents whose prudent strategy in $\text{HMIP}(M^f, \{\alpha_m\}_m)$ is performing method m .

This is a weak requirement since we only need to set sufficiently high coefficients to incentivize **two** low cost agents such that for each agent (including one of the low cost agent), she will believe there exists a low cost agent who will be incentivized to report all levels of information. potent coefficients exist since we can always set the coefficient of the highest level information sufficiently high and the coefficients of other levels arbitrarily close to zero such that agents will be incentivized to invest the highest level effort. We use our peer grading example to show how to set potent coefficients. With our example, we will see we can always set the optimal potent coefficients that minimize the mechanism's cost by solving a linear programming.

In our example, the 2 low cost agents need efforts 1,2,5 to perform m_l, m_w, m_q respectively and 8 high cost agents need efforts 1,4,10. With the above set up, we need $\alpha_q * 0.0115 + \alpha_w * 0.2259 + \alpha_l * 0.6931 - 5 > \max\{\alpha_q * 0.0041 + \alpha_w * 0.2218 + \alpha_l * 0.6931 - 2, \alpha_l * 0.6931 - 1, 0\}$ to make the coefficients potent and we also want to minimize the mechanism's cost which is

$$2 * (\alpha_q * 0.0115 + \alpha_w * 0.2259 + \alpha_l * 0.6931) + 8 * \begin{cases} v_q := \alpha_q * 0.0115 + \alpha_w * 0.2259 + \alpha_l * 0.6931 & \text{if } v_q - 10 \geq v_w - 4, v_l - 1, 0 \\ v_w := \alpha_q * 0.0041 + \alpha_w * 0.2218 + \alpha_l * 0.6931 & \text{if } v_w - 4 \geq v_q - 10, v_l - 1, 0 \\ v_l := \alpha_l * 0.6931 & \text{if } v_l - 1 \geq v_w - 4, v_q - 10, 0 \\ 0 & \text{otherwise} \end{cases}$$

After solving this linear programming, the optimal solution is around $\alpha_l, \alpha_w, \alpha_q = \epsilon^a, 0.5562, 423.8571$ and the amount of information for performing m_l, m_w, m_q are $O(\epsilon), 1.86 + O(\epsilon),$ and $5 + O(\epsilon)$ respectively. The minimal cost is $2 * (5 + O(\epsilon)) = 10 + O(\epsilon).$

^a ϵ is an arbitrarily small positive real number, we need ϵ since we want agents will be incentivized to report the length signal as well.

THEOREM 3.7. *Given a convex function f , $HMIP(MI^f, \{\alpha_m\}_m)$ is dominant truthful; moreover, when $\{\alpha_m\}_m$ are potent for $HMIP(MI^f, \{\alpha_m\}_m)$, $HMIP(MI^f, \{\alpha_m\}_m)$ is potent and dominant truthful.*

The proof of the theorem uses the information monotonicity of MI^f . The key observation in the proof is that *applying any strategy to the information is essentially data processing and thus erodes information*.

PROOF FOR THEOREM 3.7. In order to show the dominant truthful property, we will show for each agent, fixing any other agents' strategies, she can maximize her payment as well as her utility by reporting her received information honestly. The information monotonicity property of f -mutual information MI^f (Fact 2.3) says any data processing decreases the (conditional) mutual information. For each $m \in M_{-i}$, fixing the strategies other agents use, the distribution of $\hat{\Psi}_{-i}^m$, whose randomness comes from random variable Ψ_{-i}^m and the agents' strategies, is also fixed. Any strategy (data processing) agent i applies to her received signals decreases

$$MI^f(\text{her received signals}; \hat{\Psi}_{-i}^m | \{\hat{\Psi}_{-i}^{m'}\}_{m' < m, m' \in M_{-i}}).$$

Thus, for agent i , honestly reporting her received signals maximizes her payment no matter what strategies other agents use.

We start to show HMIP is potent when the coefficients are potent. When the coefficients are potent, for every agent i , when she believes everyone else plays a prudent strategy, she will believe for each m , there exists an agent $j(m) \neq i$ who reports $\{\Psi_j^\ell\}_{\ell \leq m}$ to the mechanism. Thus, $M_{-i} = M$. In this case, agent i 's optimal payment for performing method m_i will be

$$\sum_{m \in M} \alpha_m MI^f(\{\Psi_i^\ell\}_{\ell \leq m_i}; \Psi_j^m | \{\Psi_j^{m'}\}_{m' < m}) = \sum_{m \in M} \alpha_m MI^f(\{\Psi_i^\ell\}_{\ell \leq m_i}; \Psi_{-i}^m | \{\Psi_{-i}^{m'}\}_{m' < m})$$

which is the amount of information acquired with method m_i (Definition 3.3) and can be obtained by honestly reporting all received information. The optimality of the payment is due to the information monotonicity of MI^f . In this case, her optimal strategy is her prudent strategy. Therefore, HMIP is potent. $\square$

HMIP provides a framework to design information elicitation mechanisms for our hierarchical information model. To apply the HMIP framework in different settings, it remains to design the report requirement for agents and to use agents' reports to calculate the (conditional) mutual information without underlying distributions. We apply HMIP in both the multi-task setting (Section 4) and the single-task setting (Section 5).

4 MULTI-TASK SETTING

In this section, we will apply the HMIP framework in the multi-task setting where each agent receives a random batch of *a priori similar* tasks: tasks are i.i.d samples for all agents. For every agent, before she invests any effort, for each m , for all tasks, she has the same prior belief for the signals she and other agents will receive by performing m .

Prior work [6, 20, 36] also makes this assumption; however, in their setting, it is much stronger than in ours. For example, it insists that the only "signal" included in a prompt is for the correct answer. In reality, some false answers are more appealing than others (see Example 1.1 where

Kansas is an unlikely answer). In our model, these appealing false answers can be modeled as “cheap” information instead of being assumed away.

Note that in the multi-task setting, we allow agents to have different priors and only require that for every agent, her prior satisfies our assumptions.

4.1 Known Information Structure and a Small Number of tasks

In order to avoid agents “getting something for nothing” by reporting the cheap signals instead of the expensive signals (e.g. giving a high quality grade when there are no typos in Example 1.2), we reward agents the information score of expensive signals according to not only their agreements but also the diversity of their answers *conditioning on the tasks which have the same cheap signals*. (e.g. the essays which all have no typos). We will show this idea is essentially the application of HMIP framework when MI^f is chosen to be MI^{tvd} .

We assume agents’ signals are *positively correlated*. In the peer grading example, this assumption means that for every agent, receiving $\odot$ for the quality signal will increase her belief for the probability other agents receive $\odot$ for the quality signal. We defer the formal definition to the full version.

Multi-task Hierarchical Mutual Information Mechanism (Multi-HMIM($\{\alpha_m\}_m$)).

Report Each agent i is assigned a random batch of tasks (at least two). For each task t which is assigned to agent i , she is asked to report both the method $m_i(t)$ she performed on task t and method $m_i(t)$ ’s output $\psi_i^{m_i(t)}(t)$; for each $m \neq m_i(t)$, agent i is asked to optionally report her signal $\psi_i^m(t)$. We denote her actual report for her performed method and signal for every method m by $\hat{m}_i(t)$ and $\hat{\psi}_i^m(t)$ respectively.

Information Score For each method m , the mechanism collects agent i ’s method m signals and records them via a T dimensional vector $\hat{\psi}_i^m$.

$$\text{The } t^{th} \text{ coordinate of } \hat{\psi}_i^m \text{ is } \begin{cases} \hat{\psi}_i^m(t), & \text{if agent } i \text{ provides the} \\ & \text{method } m\text{'s output } \hat{\psi}_i^m(t) \\ & \text{for task } t; \\ \emptyset, & \text{otherwise} \end{cases}$$

We define $\hat{\psi}_{-i}^m$ as a vector where the t^{th} coordinate of $\hat{\psi}_{-i}^m$ is

$$\begin{cases} \hat{\psi}_{-i}^m(t), & \text{we arbitrarily pick an agent (who is not agent } i) \\ & \text{whose performed method is } \geq m \text{ for task } t \\ & \text{and provides method's } m\text{'s output for task } t; \\ & \text{we denote his report by } \hat{\psi}_{-i}^m(t); \\ \emptyset, & \text{such agent does not exist} \end{cases}$$

Agent i is paid by her information score

$$\sum_m 2\alpha_m \text{Corr}(\hat{\psi}_i^m; \hat{\psi}_{-i}^m | \{\hat{\psi}_{-i}^{m'}\}_{m' < m})$$

and $\text{Corr}(\cdot)^4$ is a random algorithm defined in Algorithm 1.

⁴ $\text{Corr}(\cdot; \cdot)$ is essentially the same concept as the payment schemes in Dasgupta and Ghosh [6], Kong and Schoenebeck [20], Shnayder et al. [36]. $\text{Corr}(\cdot; \cdot | \cdot)$ is a new concept in this paper.

Algorithm 1 Building Block *Corr*

```

1: procedure Corr( $\mathbf{v}_1; \mathbf{v}_2$ )                                ▶ e.g.  $\mathbf{v}_1 = (\odot, \emptyset, \odot, \odot, \odot)$ ,  $\mathbf{v}_2 = (\odot, \odot, \odot, \odot, \emptyset)$ 
2:   if either  $\mathbf{v}_1$  or  $\mathbf{v}_2$  has fewer than two non-empty entries then return 0
3:   else
4:      $B \subset [M] \leftarrow$  the set of entries where both  $\mathbf{v}_1$  and  $\mathbf{v}_2$  are not empty    ▶  $B \leftarrow \{1, 3, 4\}$ 
5:     if  $B = \emptyset$  then return 0
6:     else
7:       for  $t_B \in B$  do                                    ▶ We call  $t_B$  a reward task
8:          $v_1(t_1) \leftarrow$  a random non-empty entry in  $\mathbf{v}_1$ 
9:          $v_2(t_2) \leftarrow$  a random non-empty entry in  $\mathbf{v}_2$ ,  $t_2 \neq t_1$     ▶  $v_1(t_1) \leftarrow \odot$ 
10:         $\text{Corr}_{t_B} \leftarrow \mathbb{1}(v_1(t_B) = v_2(t_B)) - \mathbb{1}(v_1(t_1) = v_2(t_2))$     ▶  $v_2(t_2) \leftarrow \odot$ 
11:         $\text{Corr}_{t_B} \leftarrow 0$     ▶  $\text{Corr}_{t_B} \leftarrow 0$ 
12:   return  $\sum_{t_B \in B} \text{Corr}_{t_B}$  and “success”    ▶ Return 0

13: procedure Corr( $\mathbf{v}_1; \mathbf{v}_2 | V$ )                                ▶ e.g.  $\mathbf{v}_1 = (\odot, \odot, \odot, \odot, \odot)$ ,  $\mathbf{v}_2 = (\odot, \odot, \odot, \odot, \odot)$ ,  $V = \{v\}$ ,
    $v = (\odot, \odot, \odot, \odot, \odot)$ 
14:    $C \leftarrow$  the set of entries where every  $v \in V$  is not empty    ▶  $C \leftarrow \{1, 2, 3, 4, 5\}$ 
15:   if  $C = \emptyset$  then return Corr( $\mathbf{v}_1; \mathbf{v}_2$ )
16:   else
17:      $t_C^* \leftarrow$  a random element in  $C$     ▶  $t_C^* \leftarrow 2$ 
18:      $D \leftarrow \emptyset$ 
19:     for  $t \in [T]$  do
20:       if for every  $v \in V$ ,  $v(t) = v(t_C^*)$  then
21:         put  $t$  in  $D$     ▶  $D = \{1, 2, 4\}$ ,  $\mathbf{v}_1(D) = \mathbf{v}_2(D) = (\odot, \odot, \odot)$ 
22:   return Corr( $\mathbf{v}_1(D); \mathbf{v}_2(D)$ )
23:   ▶ Return Corr( $\mathbf{v}_1(D); \mathbf{v}_2(D)$ ) = 0 and “success”

```

We design $\text{Corr}(\hat{\psi}_i^m; \hat{\psi}_{-i}^m | \{\hat{\psi}_{-i}^{m'}\}_{m' < m})$ to be an unbiased estimator of $MI^{tvd}(\hat{\Psi}_i^m; \hat{\Psi}_{-i}^m | \{\hat{\Psi}_{-i}^{m'}\}_{m' < m})$ ⁵ if $\hat{\Psi}_i^m$ and $\hat{\Psi}_{-i}^m$ are *positively correlated*. Thus, in Multi-HMIM, agents are essentially paid based on the (conditional) mutual information by picking a special f -mutual information— MI^{tvd} , if agents are honest since we have assumed that agents’ honest signals are positively correlated. This makes our Multi-HMIM a special case of HMIP framework.

With the positively correlated and other additional natural assumptions (see full version), with similar definitions for the amount of information and potent coefficients with HMIP, we can show Multi-HMIM is potent as well by mapping it into the HMIP framework. Moreover, we define truthful strategy in Multi-HMIM as the strategy where the agent honestly reports her performed method and that method’s output for each task; for other methods’ outputs, either she chooses to not report or she reports honestly. We can show Multi-HMIM is truthful: for each agent, when she believes other agents play a truthful strategy, she can maximize her expected utility by playing a truthful strategy as well. We defer the formal analysis of Multi-HMIM to the full version.

4.2 Learning Information Structure with a Large Number Tasks

ASSUMPTION 4.1 (δ_0 -GAP). *For each m , we assume that for every $i \neq j$, each $m' \neq m$*

$$MI^f(\Psi_i^m; \Psi_j^m) > \frac{1}{\delta_0} \quad MI^f(\Psi_i^m; \Psi_j^{m'}) < \frac{1}{\delta_0}$$

⁵In the current paper, $\hat{\psi}_i^m$ means vector, $\hat{\Psi}_i^m$ means random variable.

The above assumption guarantees that when we can accurately learn the f -mutual information between two agents' answer vectors, we can accurately classify the answer vectors and then learn the maximal method's outputs correctly.

Learning-based Multi-HMIM($\mathcal{L}$).

Report Each agent i is assigned T tasks and asked to perform the same method for all tasks. For agent i who performs method m_i , she is asked to report her own answer vector

$$\psi_i^{m_i} = (\psi_i^{m_i}(1), \psi_i^{m_i}(2), \dots, \psi_i^{m_i}(T))$$

and, for each method $m \neq m_i$, is asked to optionally report her answer vector ψ_i^m . We denote the set of methods whose outputs are reported by agent i as M_i and the actual answer vector she reports for each method $\ell \in M_i$ as $\hat{\psi}_i^\ell$. Agent i can name the methods freely⁶.

Learning Information Structure We define the distance between $\hat{\psi}_i^m$ and $\hat{\psi}_j^{m'}$ as $\frac{1}{MI^f(\hat{\psi}_i^m, \hat{\psi}_j^{m'})}$.

The mechanism starts to cluster answer vectors. A set of answer vectors are clustered into one cluster if and only if their pairwise distance is less than δ_0 . A cluster may have ≥ 1 answer vector(s).

For two clusters m_1, m_2 , $m_1 > m_2$ if and only if there exists an agent who's own answer vector is in cluster m_1 and also provides an answer vector which is classified in cluster m_2 . The mechanism picks positive real values for the type payment scale α_m according to a rule $\mathcal{L}$.

Information Score The mechanism learns the information structure using all agents' reports excluding agent i . We denote the set of clusters by M_{-i} . For each cluster $m \in M_{-i}$, the mechanism randomly picks an answer vector, denoted $\hat{\psi}_{-i}^m$, from it.

Agent i is paid her information score:

$$\sum_{m \in M_{-i}} \alpha_m MI^f(\{\hat{\psi}_i^\ell\}_{\ell \in M_i}, \hat{\psi}_{-i}^m | \{\hat{\psi}_{-i}^{m'}\}_{m' < m, m' \in M_{-i}})$$

which can be calculated accurately when the number of tasks is large.

We define $\alpha(\mathcal{L}) := \{\alpha_m(\mathcal{L})\}_m$ as the coefficients determined by $\mathcal{L}$, given that the mechanism has access to all levels of honest answer vectors. Here the amount of information and prudent strategy are defined by the same way in HMIP, except that the coefficients are $\alpha(\mathcal{L})$.

We also define potent $\mathcal{L}$ such that $\alpha(\mathcal{L})$ is potent in the definition in HMIP. We can show with the a priori similar and random order assumption, Learning-based multi-HMIM is dominant truthful. With the gap assumption, when the rule $\mathcal{L}$ is potent, Learning-based multi-HMIM is potent, dominant truthful and will output the hierarchical information structure as well as the maximal level(s) answer vector given that agents play prudent strategy.

Learning-based multi-HMIM can be mapped to HMIP since a large number of tasks and the gap assumption allow the estimation of the prior and the learning of the information structure and correct clusters. Note that even if the mechanism clusters incorrectly, the mechanism is still dominant truthful since even each agent is paid by the mutual information between her information and "wrong" information, the information monotonicity still incentivizes the agent to report all information she has. Thus we do not need the gap assumption for the dominant truthfulness. With the gap assumption, we can cluster agents correctly and use Theorem 3.7 to show the potent property. We defer the formal statement and proof to the full version. Moreover, we want to emphasize that our mechanisms work even if every agent only has a piece of correct information for the information structure.

⁶The mechanism will ignore the name of the methods and only record the relationship that the other answer vectors reported by agent i have lower level than agent i 's own answer vector.

5 SINGLE-TASK SETTING

In the single task setting without known prior, the literature usually assumes a common prior assumption and follows a signal-prediction framework [28]—asking agents not only her signal but also her prediction. To achieve truthfulness for ≥ 3 agents, Radanovic and Faltings [31] and Kong and Schoenebeck [21] punish each agent if her predictions differ from the predictions of other agents who report the same signals with her, and reward each agent for the accuracy of her prediction. Therefore, for each agent, to maximize her accuracy reward, she will honestly report her predictions. To avoid the punishment for the “inconsistency”, she will honestly report her received signals as well because of the following commonly assumed assumption—*common prior and stochastic relevance*: for every agent i and j , they will have the same belief for the distribution of the signals received by other agents if and only if they receive the same signals.

5.1 Applying HMIP in the Single-task Setting

We follow the previous “signal-prediction” framework and “punishing inconsistency” idea in the hierarchical information case. We ask agents their received signals and predictions for different levels. We pay each agent the accuracy of her forecasts. The high expertise agents have accurate predictions for even high cost information reports while the low expertise agents only have accurate prediction for low cost information. Therefore, high expertise agents will be paid more.

Single-HMIM($PS, \{\alpha_m\}_m$).

Report (signals, predictions) Each agent i who performs method m_i is asked to report her received signals $\{\sigma_i^m\}_{m \leq m_i}$ and her forecast $p_i^{m_i}$ for $\Psi_{-i}^{m_i}$. For each $m \neq m_i$, she is asked to optionally report her forecast p_i^m for Ψ_{-i}^m . We denote her report for her received signals as $\{\hat{\sigma}_i^m\}_{m \leq \hat{m}_i}$ and her prediction report as $\{\hat{p}_i^m\}_{m \in M_i}$ where M_i is the set of methods whose outputs are predicted by agent i .

Prediction Score We define M_{-i} as the set of methods whose outputs are reported by an agent who is not agent i . For each $m \in M_{-i}$, we pick an arbitrary reference agent $j \neq i$ whose performed method is higher than m and denote his report for method m 's output by $\hat{\sigma}^m$. Agent i 's prediction score is $\sum_{m \in M_{-i} \cap M_i} \alpha_m PS(\sigma_i^m, \hat{p}_i^m)$.

Information Score If there is no other agent who reports the same signals as agent i , then agent i 's information score is 0. Otherwise, arbitrarily pick a reference agent $j \neq i$ from the agents who report the same signals as agent i . Agent i 's information score is minus the inconsistency between her prediction report and agent j 's prediction report, that is,

$$-\left(\sum_{m \in M_i \cap M_j} \alpha_m (PS(\hat{p}_j^m, \hat{p}_j^m) - PS(\hat{p}_j^m, \hat{p}_i^m)) \right).$$

In Single-HMIM, the payment of each agent is $\alpha * \text{Information Score} + \beta * \text{Prediction Score}$.

Definition 5.1 (Truthful strategy in Single-HMIM). For each agent i who performed method m_i , she plays truthful strategy if she honestly report her received signals $\{\sigma_i^m\}_{m \leq m_i}$ and forecast for $\Psi_{-i}^{m_i}$ and for each $m \neq m_i$, either she does not report or she reports her forecast for Ψ_{-i}^m honestly.

With proper definition for the amount of information and potent coefficients, we will show with the common prior assumption, single-HMIM is strictly truthful; moreover, when the coefficients are potent for single-HMIM, single-HMIM is potent and strictly truthful.

The strictly truthful property follows from the common prior and stochastic relevance assumption. The potent property follows from the definition of prudent strategy and potent coefficients. We defer the formal proof to the full version.

6 DISCUSSION AND FUTURE WORK

Although our work is theoretical, it is not far from applications. For example, in some situations, we could simplify the information structure by roughly dividing it into two levels where the higher level requires more time. Then we can apply our Multi-HMIM or Single-HMIM to peer grading or any other situations where certain agents are only given 15 seconds to grade a work and others are expected to do a good job. We value the information *conditioning* on the reports provided by the “15 seconds” agents. We could also use machine learning to obtain the lower level information.

Race, sex, and other stereotypes can be used as “cheap” signals. One possible future application of this work is to address fairness in crowdsourcing.

There are several limitations of our mechanisms. Our Multi-HMIM requires that the hierarchy structure of the information is common knowledge. Our key assumption, that more sophisticated agents know the information of less sophisticated agents may not always hold, especially if the less sophisticated agents collude to report seemingly irrelevant information (e.g., the third letter of the third word).

Our analysis of the Learning-based HMIM requires the agents to perform an infinite number of tasks. One future direction is having the sample complexity analysis. Our Single-HMIM still requires a forecast report. A future direction is making the forecast report optional in some scenarios (we might ask agents for the most common answer if they believe it is not their reported answer). Another future direction is the information cost elicitation: tuning the coefficients of the mechanisms such that the payment matches the actual effort required by agents.

A key future direction is to test our mechanism and explore the limitations in practice by performing real-world experiments. This would require designing a suitable user interface.

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, EC '17, Cambridge, MA, USA, June 26-30, 2017*. 81–98.
- [2] Syed Mumtaz Ali and Samuel D Silvey. 1966. A general class of coefficients of divergence of one distribution from another. *Journal of the Royal Statistical Society. Series B (Methodological)* (1966), 131–142.
- [3] Y. Chen, D.F. Bacon, I. Kash, D.C. Parkes, M. Rao, and M. Sridharan. 2012. Predicting Your Own Effort. In *Proceedings of the 11th International Conference on Autonomous and Multiagent Systems (AAMAS 2012)*.
- [4] Thomas M Cover and Joy A Thomas. 2006. Elements of information theory 2nd edition. (2006).
- [5] Imre Csiszár, Paul C Shields, et al. 2004. Information theory and statistics: A tutorial. *Foundations and Trends® in Communications and Information Theory* 1, 4 (2004), 417–528.
- [6] Anirban Dasgupta and Arpita Ghosh. 2013. Crowdsourced judgement elicitation with endogenous proficiency. In *Proceedings of the 22nd international conference on World Wide Web*. International World Wide Web Conferences Steering Committee, 319–330.
- [7] 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.
- [8] Boi Faltings, Radu Jurca, Pearl Pu, and Bao Duy Tran. 2014. Incentives to counter bias in human computation. In *Second AAAI Conference on Human Computation and Crowdsourcing*.
- [9] A. Gao, J. R. Wright, and K. Leyton-Brown. 2016. Incentivizing Evaluation via Limited Access to Ground Truth: Peer-Prediction Makes Things Worse. *ArXiv e-prints* (June 2016). arXiv:cs.GT/1606.07042
- [10] Arpita Ghosh, Satyen Kale, and Preston McAfee. 2011. Who moderates the moderators?: crowdsourcing abuse detection in user-generated content. In *Proceedings of the 12th ACM conference on Electronic commerce*. ACM, 167–176.
- [11] Tilmann Gneiting and Adrian E Raftery. 2007. Strictly proper scoring rules, prediction, and estimation. *J. Amer. Statist. Assoc.* 102, 477 (2007), 359–378.
- [12] Sharad Goel, Daniel M. Reeves, and David M. Pennock. 2009. Collective revelation: A mechanism for self-verified, weighted, and truthful predictions. In *Proceedings of the 10th ACM conference on Electronic commerce (EC 2009)*.
- [13] Radu Jurca and Boi Faltings. 2006. Minimum payments that reward honest reputation feedback. In *Proceedings of the 7th ACM conference on Electronic commerce (EC 2006)*.

- [14] Radu Jurca and Boi Faltings. 2007. Robust Incentive-Compatible Feedback Payments. In *Trust, Reputation and Security: Theories and Practice*, Vol. 4452. Springer-Verlag, 204–218.
- [15] Radu Jurca and Boi Faltings. 2008. Incentives for expressing opinions in online polls. In *Proceedings of the 9th ACM conference on Electronic commerce (EC 2008)*.
- [16] Radu Jurca and Boi Faltings. 2009. Mechanisms for making crowds truthful. *J. Artif. Int. Res.* 34, 1 (March 2009).
- [17] R. Jurca and B. Faltings. 2011. Incentives for Answering Hypothetical Questions. In *Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC 2011)*. ACM.
- [18] Vijay Kamble, Nihar Shah, David Marn, Abhay Parekh, and Kannan Ramachandran. 2015. Truth Serums for Massively Crowdsourced Evaluation Tasks. *arXiv preprint arXiv:1507.07045* (2015).
- [19] 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.
- [20] Y. Kong and G. Schoenebeck. 2016. A Framework For Designing Information Elicitation Mechanisms That Reward Truth-telling. *ArXiv e-prints* (May 2016). arXiv:cs.GT/1605.01021
- [21] Yuqing Kong and Grant Schoenebeck. 2018. Equilibrium selection in information elicitation without verification via information monotonicity. In *9th Innovations in Theoretical Computer Science, ITCS 2018*, Vol. 94.
- [22] N. Lambert and Y. Shoham. 2008. Truthful surveys. *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE 2008)* (2008).
- [23] Yang Liu and Yiling Chen. 2016. Learning to incentivize: eliciting effort via output agreement. *arXiv preprint arXiv:1604.04928* (2016).
- [24] Yang Liu and Yiling Chen. 2016. Sequential Peer Prediction: Learning to Elicit Effort using Posted Prices. *arXiv preprint arXiv:1611.09219* (2016).
- [25] Debmalya Mandal, Matthew Leifer, David C Parkes, Galen Pickard, and Victor Shnayder. 2016. Peer Prediction with Heterogeneous Tasks. *arXiv preprint arXiv:1612.00928* (2016).
- [26] John McCoy and Drazen Prelec. 2017. A statistical model for aggregating judgments by incorporating peer predictions. *arXiv preprint arXiv:1703.04778* (2017).
- [27] N. Miller, P. Resnick, and R. Zeckhauser. 2005. Eliciting informative feedback: The peer-prediction method. *Management Science* (2005), 1359–1373.
- [28] D. Prelec. 2004. A Bayesian Truth Serum for subjective data. *Science* 306, 5695 (2004), 462–466.
- [29] Drazen Prelec, H Sebastian Seung, and John McCoy. 2017. A solution to the single-question crowd wisdom problem. *Nature* 541, 7638 (2017), 532–535.
- [30] 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, AAAI 2013*. 833–839.
- [31] Goran Radanovic and Boi Faltings. 2014. Incentives for truthful information elicitation of continuous signals. In *Twenty-Eighth AAAI Conference on Artificial Intelligence*.
- [32] Goran Radanovic and Boi Faltings. 2015. Incentive schemes for participatory sensing. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*. 1081–1089.
- [33] Goran Radanovic and Boi Faltings. 2015. Incentives for Subjective Evaluations with Private Beliefs. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- [34] Blake Riley. 2014. Minimum truth serums with optional predictions. In *Proceedings of the 4th Workshop on Social Computing and User Generated Content (SC14)*.
- [35] Blake Riley. 2015. Mechanisms for Making Accurate Decisions in Biased Crowds. (2015).
- [36] 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.
- [37] Peter Welinder, Steve Branson, Serge J Belongie, and Pietro Perona. 2010. The Multidimensional Wisdom of Crowds.. In *NIPS*, Vol. 23. 2424–2432.
- [38] Robert L Winkler. 1969. Scoring rules and the evaluation of probability assessors. *J. Amer. Statist. Assoc.* 64, 327 (1969), 1073–1078.
- [39] Jens Witkowski, Bernhard Nebel, and David C Parkes. 2014. *Robust Peer Prediction Mechanisms*. Ph.D. Dissertation. Ph. D. Dissertation, Albert-Ludwigs-Universitat Freiburg: Institut fur Informatik.
- [40] Jens Witkowski and David C Parkes. 2012. Peer prediction without a common prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce*. ACM, 964–981.
- [41] Jens Witkowski and David C Parkes. 2013. Learning the prior in minimal peer prediction. In *Proceedings of the 3rd Workshop on Social Computing and User Generated Content at the ACM Conference on Electronic Commerce*. Citeseer, 14.
- [42] Peter Zhang and Yiling Chen. 2014. Elicitability and knowledge-free elicitation with peer prediction. In *Proceedings of the 2014 international conference on Autonomous agents and multi-agent systems*. 245–252.
- [43] Denny Zhou, Sumit Basu, Yi Mao, and John C Platt. 2012. Learning from the wisdom of crowds by minimax entropy. In *Advances in Neural Information Processing Systems*. 2195–2203.