

Crowdsourced Judgement Elicitation with Endogenous Proficiency

Anirban Dasgupta
Yahoo! Labs
Sunnyvale, CA, USA
anirban@yahoo-inc.com

Arpita Ghosh
Cornell University
Ithaca, NY, USA
arpitaghosh@cornell.edu

ABSTRACT

Crowdsourcing is now widely used to replace judgement or evaluation by an expert authority with an aggregate evaluation from a number of non-experts, in applications ranging from rating and categorizing online content all the way to evaluation of student assignments in massively open online courses (MOOCs) via peer grading. A key issue in these settings, where direct monitoring of both effort and accuracy is infeasible, is incentivizing agents in the ‘crowd’ to put in *effort* to make good evaluations, as well as to truthfully report their evaluations. We study the design of mechanisms for crowdsourced judgement elicitation when workers strategically choose both their reports and the effort they put into their evaluations. This leads to a new family of information elicitation problems with unobservable ground truth, where an agent’s proficiency—the probability with which she correctly evaluates the underlying ground truth—is *endogenously* determined by her strategic choice of how much effort to put into the task.

Our main contribution is a simple, new, mechanism for binary information elicitation for multiple tasks when agents have *endogenous proficiencies*, with the following properties: (i) Exerting maximum effort followed by truthful reporting of observations is a Nash equilibrium. (ii) This is the equilibrium with *maximum payoff* to all agents, even when agents have different maximum proficiencies, can use mixed strategies, and can choose a different strategy for each of their tasks. Our information elicitation mechanism requires only minimal bounds on the priors, asks agents to only report their own evaluations, and does not require any conditions on a diverging number of agent reports per task to achieve its incentive properties. The main idea behind our mechanism is to use the presence of multiple tasks and ratings to estimate a reporting statistic to identify and penalize low-effort agreement—the mechanism rewards agents for agreeing with another ‘reference’ report on the same task, but also penalizes for *blind agreement* by subtracting out this statistic term, designed so that agents obtain rewards only when they put in effort into their observations.

Categories and Subject Descriptors

J.4 [Computer Applications]: Social & Behavioral Sciences

Keywords

Crowdsourcing, Information elicitation, Peer prediction, Mechanism design, Game theory

Copyright is held by the International World Wide Web Conference Committee (IW3C2). IW3C2 reserves the right to provide a hyperlink to the author’s site if the Material is used in electronic media. *WWW 2013*, May 13–17, 2013, Rio de Janeiro, Brazil. ACM 978-1-4503-2035-1/13/05.

1. INTRODUCTION

Crowdsourcing, where a problem or task is broadcast to a crowd of potential participants for solution, is used for an increasingly wide variety of tasks on the Web. One particularly common application of crowdsourcing is in the context of making evaluations, or judgements—when the number of evaluations required is too large for a single expert, a solution is to replace the expert by an evaluation aggregated from a ‘crowd’ recruited on an online crowdsourcing platform such as Amazon Mechanical Turk. Crowdsourced judgement elicitation is now used for a number of applications such as image classification and labeling, judging the quality of online content, identifying abusive or adult content, and most recently for peer grading in online education, where Massively Open Online Courses (MOOCs) with enrollment in the hundreds of thousands crowdsource the problem of evaluating assignments submitted by students back to the class itself. While one issue in the context of crowdsourcing evaluations is how best to aggregate the evaluations obtained from the crowd, there is also a key question of *eliciting the best possible evaluations* from the crowd in the first place.

The problem of designing incentive mechanisms for such crowdsourced judgement elicitation scenarios has two aspects. First, suppose each worker has already evaluated, or formed a judgement on, the tasks allocated to her. Since the ‘ground truth’ for each task is unknown to the system, a natural solution is to reward workers based on other workers’ reports for the same task (this being the only available source of information about this ground truth)¹. The problem of designing rewards to incentivize agents to truthfully report their observation, rather than, for example, a report that is more likely to agree with other agents’ reports, is an *information elicitation* problem with unobservable ground truth. Information elicitation has been recently been addressed in the literature in the context of eliciting opinions online (such as user opinions about products, or experiences with service providers); see §1.1. However, in those settings, agents (users) have *already* formed an opinion after receiving their signal (for example a user who buys a product forms an opinion about it after buying it)—so agents only need to be incentivized to incur the cost to report this opinion and find it more profitable to report their opinions *truthfully* than to report a different opinion.

In the crowdsourcing settings we consider, however, the user does not have such a pre-formed, or experiential, opinion anyway,

¹It is of course infeasible for a requester to monitor every worker’s performance on her task, since this would be a problem of the same scale as simply performing all the tasks herself. We also note that a naive approach of randomly checking some subset of evaluations, either via inserting tasks with known responses, or via random checking by the requester, turns out to be very wasteful of effort at the scale necessary to achieve the right incentives.

but rather forms a judgement *as part of her task*— further, the accuracy of this judgement depends on whether or not the agent puts in effort into it (for instance, a worker evaluating whether images contain objectionable content could put in no effort and declare all images to be clean, or put in effort into identifying which images are actually appropriate; a similar choice applies in other contexts like peer-grading). A key issue in these crowdsourced judgement elicitation scenarios is therefore *incentivizing effort*²— that is, ensuring that agents *make* the best judgements that they possibly can (in addition, of course, to ensuring that they then truthfully report this observation). This leads to a new kind of information elicitation problem where an agent’s proficiency now *depends* on her effort choice, and so is *endogenous* and unknown to the system— even if an agent’s maximum proficiency is known, the actual proficiency with which she performs a task is an endogenous, strategic choice and therefore cannot be assumed as fixed or given.

A mechanism for information elicitation in this setting should make it ‘most beneficial’, if not the only beneficial strategy, for agents to not just *report* their observations truthfully, but to also *make* the best observations they can in the first place. Also, it is even more important now to ensure that the payoffs from all agents always blindly reporting the same observation (for instance, declaring all content to be good) are strictly smaller than the payoffs from truthfully reporting what was actually observed, since declaring all tasks to be of some predecided type requires no effort and therefore incurs no cost, whereas actually putting in effort into making observations will incur a nonzero cost. Finally, unlike mechanisms designed for settings where a large audience is being polled for its opinion about a single event, a mechanism here must retain its incentive properties even when there are only a few reports per task— this is because it can be infeasible, due to either monetary or effort constraints, to solicit reports from a large number of agents for each task. (For example, the number of tasks in peer grading scales *linearly* with the number of agents, limiting the number of reports available for each task since each student can only grade a few assignments; similarly, the total cost to the requester in crowdsourcing platforms such as Amazon Mechanical Turk scales linearly with the number of workers reporting on each task). How can we elicit the best possible evaluations from agents whose proficiency of evaluation depends on their strategically chosen effort, when the ground truth as well as the effort levels of agents are unobservable to the mechanism?

Our Contributions. We introduce a model for information elicitation with *endogenous proficiency*, where an agent’s strategic choice of whether or not to put in effort into a task endogenously determines her proficiency (the probability of correctly evaluating the ground truth) for that task. We focus on the design of mechanisms for binary information elicitation, *i.e.*, when the underlying ground truth is binary (corresponding to eliciting ‘good’ or ‘bad’ ratings). While generalizing to an arbitrary underlying type space is an immediate direction for further work, we note that a number of interesting judgement and evaluation tasks, for example identifying adult content or correctness evaluation, are indeed binary; also, even very recent literature providing improved mechanisms for information elicitation (e.g. [17, 19]), as well as experimental

²We thank David Evans (VP Education, Udacity) for pointing out this issue in the context of peer-grading applications— while students might put in their best efforts on grading screening assignments to ensure they demonstrate the minimum proficiency required to be allowed to grade, how can we be sure that they will continue to work with the same proficiency when grading homeworks outside of this screening set?

work on the performance of elicitation mechanisms [16, 5], focuses on models with binary ground truth.

Our main contribution is a simple, new, mechanism for binary information elicitation for multiple tasks when agents have endogenous proficiencies. Our mechanism has the following incentive properties.

- (i) Exerting maximum effort followed by truthful reporting of observations is a Nash equilibrium.
- (ii) This is the equilibrium with *maximum payoff* to all agents, *even* when agents have different maximum proficiencies, can use *mixed* strategies, and can choose a different strategy for each of their tasks.

Showing that full-effort truthtelling leads to the maximum reward amongst all equilibria (including those involving mixed strategies) requires arguing about the rewards to agents in all possible equilibria that may arise. To do this, we use a matrix representation of strategies where every strategy can be written as a convex combination of ‘basis’ strategies, so that maximizing a function over the set of all possible strategies is equivalent to a maximization over the space of coefficients in this convex combination. This representation lets us show that the reward to an agent over all possible strategy choices (by herself and other agents), and therefore over all equilibria, is maximized when all agents use the strategy of full-effort truthful reporting.
- (iii) Suppose there is some positive probability, however small, that there is some ‘trusted’ agent for each task who will report on that task truthfully with proficiency greater than half. Then the equilibrium where all agents put in full effort and report truthfully on all their tasks is essentially the only equilibrium of our mechanism, *even if* the mechanism does not know the identity of the trusted agents.

We note that our mechanism requires only minimal bounds on the priors and imposes no conditions on a diverging number of agent reports per task to achieve its incentive properties— to the best of our knowledge, previous mechanisms for information elicitation do not provide all these guarantees simultaneously, even when proficiency is not an endogenously determined choice (see §1.1 for a discussion).

The main idea behind our mechanism \mathcal{M} is the following. With just one task, it is difficult to distinguish between agreement rising from high-effort observations of the same ground truth, and ‘blind’ agreement achieved by the low-effort strategy of always making the same report. We use the presence of *multiple* tasks and ratings to distinguish between these two scenarios and appropriately reward or penalize agents to incentivize high effort— \mathcal{M} rewards an agent i for her report on task j for agreeing with another ‘reference’ agent $r_j(i)$ ’s report on the same task, but also *penalizes for blind agreement* by subtracting out a statistic term corresponding to the part of i and $r_j(i)$ ’s agreement on task j that is to be ‘expected anyway’ given their reporting statistics estimated from other tasks. This statistic term is chosen so that there is no benefit to making reports that are independent of the ground truth; the incentive properties of the mechanism follow from this property that agents obtain positive rewards *only* when they put effort into their evaluations.

1.1 Related Work

The problem of designing incentives for crowdsourced judgement elicitation is closely related to the growing literature on information elicitation mechanisms. The key difference between this

literature (discussed in greater detail below) and our work is that agents in the settings motivating past work have opinions that are *experientially formed* anyway— independent, and outside of, any mechanisms to elicit opinions— so that agents only need be incentivized to participate and truthfully report these opinions. In contrast, agents in the crowdsourcing settings we study *do not have such experientially formed opinions* to report— an agent makes a judgement only because it is part of her task, expending effort to *form* her judgement, and therefore must be incentivized to both expend this effort and then to truthfully report her evaluation. There are also other differences in terms of the models and guarantees in previous mechanisms for information elicitation; we discuss this literature below.

The peer-prediction method, introduced by Miller, Resnick and Zeckhauser [12], is a mechanism for the information elicitation problem for general outcome spaces where truthful reporting is a Nash equilibrium, uses proper scoring rules to reward agents for reports that are predictive of other agents’ reports. The main difference between our mechanism and [12], as well as other mechanisms based on the peer prediction method [8, 9, 18, 10, 19], is in the model of agent proficiency. In peer-prediction models, while agents can decide whether to incur the cost to *participate* (*i.e.*, submit their opinion), an agent’s proficiency— the distribution of opinions or evaluations conditional on ground truth— is *exogenously* determined (and common to all agents, and in most models, known to the center). That is, an agent might not participate at all, but if she does participate she is assumed to have some known proficiency. In contrast, in our setting, agents can choose not just whether or not to participate, but also *endogenously* determine their proficiency conditional on participating through their effort choice. Thus, while peer prediction mechanisms do need to incentivize agents to participate (by submitting a report), they then know the proficiency of agents who do submit a report, and therefore can, and do, dispense rewards that crucially use knowledge of this proficiency. In contrast, even agents who do submit reports in our setting cannot be assumed to be using their maximum proficiency to make their evaluations, and therefore cannot be rewarded based on any assumed level of proficiency. Additionally, truthtelling, while an equilibrium, is *not necessarily the maximum-reward equilibrium* in these existing peer-prediction mechanisms— [6] shows that for the mechanisms in [12, 8], the strategies of always reporting ‘good’ or always reporting ‘bad’ both constitute Nash equilibria, at least one of which generates *higher* payoff than truthtelling. Such blind strategy equilibria can be eliminated and honest reporting made the unique Nash equilibrium by designing the payments as in [10], but this needs agents to be restricted to pure reporting strategies, and requires full knowledge of the prior and conditional probability distributions to compute the rewards.

The Bayesian Truth Serum (BTS) [15] is another mechanism for information elicitation with unobservable ground truth. BTS does not use the knowledge of a common prior to compute rewards, but rather collects two reports from each agent— an ‘information’ report which is the agent’s own observation, as well as a ‘prediction’ report which is the agent’s prediction about the distribution of information reports from the population— and uses these to compute rewards such that truthful reporting is the highest-reward Nash equilibrium of the BTS mechanism. In addition to the key difference of exogenous versus endogenous proficiencies discussed above, an important limitation of BTS in the crowdsourcing setting is that it requires the number of agents n reporting on a task to diverge to ensure its incentive properties. This $n \rightarrow \infty$ requirement is infeasible in our setting due to the scaling of cost with number of reports as discussed in the introduction. [17] provides a robust BTS

mechanism (RBTS) that works even for small populations (again in the same non-endogenous proficiency model as BTS and peer prediction mechanisms), and also ensures payments are positive, making the mechanism ex-post individually rational in contrast to BTS. However, the RBTS mechanism does not retain the property of truthtelling being the highest reward Nash equilibrium— indeed, the ‘blind agreement’ equilibrium via constant reports achieves the maximum possible reward in RBTS, whereas truthtelling might in fact lead to lower rewards.

There is also work on information elicitation in conducting surveys and online polling [11, 7], both of which are not quite appropriate for our crowdsourcing setting. The mechanism in [11] is weakly incentive compatible (agents are indifferent between lying and truthtelling), while [7] presents an online mechanism that is not incentive compatible in the sense that we use and potentially requires a large (constant) number of agents to converge to the true result. For other work on information elicitation, albeit in settings very different from ours, see [3, 2, 14].

We note also that we model settings where there is indeed a notion of a ground truth, albeit unobservable, so that proficient agents who put in effort are more likely than not to correctly observe this ground truth. Peer-prediction methods, as well as the Bayesian truth serum, are designed for settings where there may be no underlying ground truth at all, and the mechanism only seeks to elicit agents’ true observations (whatever they are) which means that some agents might well be in the minority even when they truthfully report their observation— this makes the peer prediction setting ‘harder’ along the dimension of inducing truthful reports, but easier along the dimension of not needing to incentivize agents to choose to exert effort to make high-proficiency observations.

We note that our problem can also be cast as a version of a principal-agent problem with a very large number of agents, although the principal cannot directly observe an agent’s ‘output’ as in standard models. While there is a vast literature in economics on the principal-agent problem too large to describe here (see, eg, [1] and references therein), none of this literature, to the best of our knowledge, addresses our problem. Finally, there is also a large orthogonal body of work on the problem of *learning* unknown (but *exogenous*) agent proficiencies, as well as on the problem of optimally combining reports from agents with differing proficiencies to come up with the best aggregate evaluation in various models and settings. These problems of learning exogenous agent proficiencies and optimally aggregating agent reports are orthogonal to our problem of providing incentives to agents with endogenous, effort-dependent proficiencies to elicit the best possible evaluations from them.

2. MODEL

We now present a simple abstraction of the problem of designing mechanisms for crowdsourced judgement elicitation settings where agents’ proficiencies are determined by strategic effort choice.

Tasks. There are m tasks, or objects, $j = 1, \dots, m$, where each task has some underlying ‘true quality’, or type, \bar{X}_j . This true type \bar{X}_j is unknown to the system. We assume that the types are binary-valued: \bar{X}_j is either H (or 1, corresponding to high-quality) or L (or 0, for low quality) for all j ; we use 1 and H (resp. 0 and L) interchangeably throughout for convenience. The prior probabilities of H and L for all tasks are denoted by $\mathcal{P}[H]$ and $\mathcal{P}[L]$. We assume throughout that $\max(\mathcal{P}[H], \mathcal{P}[L]) < 1$, *i.e.*, that there is at least some uncertainty in the underlying qualities of the objects.

Agents. There are n workers or agents $i = 1, \dots, n$ who noisily evaluate, or form judgements on, the qualities of objects. We say agent i performs task j if i evaluates object j . Agent i ’s judgement

on task j is denoted by $\hat{X}_{ij} \in \{0, 1\}$, where \hat{X}_{ij} is 0 if i evaluates j to be of type L and \hat{X}_{ij} is 1 if i evaluates it to be H . Having made an evaluation \hat{X}_{ij} , an agent can choose to report any value $X_{ij} \in \{0, 1\}$ either based on, or independent of, her actual evaluation \hat{X}_{ij} .

We denote the set of tasks performed by an agent i by $J(i)$, and let $I(j)$ denote the set of agents who perform task j . We will assume for notational simplicity that $|J(i)| = D$ and $|I(j)| = T$ for all agents i and tasks j .

Proficiency. An agent’s *proficiency* at a task is the probability with which she correctly evaluates its true type or quality. We assume that an agent’s proficiency is an increasing function of the *effort* she puts into making her evaluation. Let e_{ij} denote agent i ’s effort level for task j : we assume for simplicity that effort is binary-valued, $e_{ij} \in \{0, 1\}$. Putting in 0 effort has cost $c_{ij}(0) = 0$, whereas putting in full effort has cost $c_{ij}(1) \geq 0$ (we note that our results also extend to a linear model with continuous effort where $e_{ij} \in [0, 1]$ and the probability $p_i(e_{ij})$ of correctly observing \bar{X}_j as well as the cost $c_i(e_{ij})$ increase linearly with e_{ij}).

An agent who puts in zero effort makes evaluations with proficiency $p_{ij}(0) = 1/2$ and does no better than random guessing, *i.e.*, $\Pr(\hat{X}_{ij} = \bar{X}_j | e_{ij} = 0) = 1/2$. An agent who puts in full effort $e_{ij} = 1$ attains her *maximum proficiency*, $\Pr(\hat{X}_{ij} = \bar{X}_j | e_{ij} = 1) = p_{ij}(1) = p_i$. Note that this maximum proficiency p_i can be different for individual agents modeling their different abilities, and need *not* be known to the center. We assume that the maximum proficiency $p_i \geq \frac{1}{2}$ for all i —this minimum requirement on agent ability can be ensured in online crowdsourcing settings by *prescreening* workers on a representative set of tasks (Amazon Mechanical Turk, for instance, offers the ability to prescreen workers [13, 4], whereas in peer-grading applications such as on Coursera, students are given a set of pre-graded assignments to measure their grading abilities prior to grading their peers, the results of which can be used as a prescreen.)

We note that our results also extend easily to the case where the maximum proficiency of an agent *depends* on whether the object is of type H or L , *i.e.*, the probabilities of correctly observing the ground truth when putting in full effort are different for different ground truths, $\Pr(\hat{X}_{ij} = \bar{X}_j | \bar{X}_j = H) \neq \Pr(\hat{X}_{ij} = \bar{X}_j | \bar{X}_j = L)$ (of course, different agents can continue to have different maximum proficiencies).

Strategies. Agents strategically choose *both* their effort levels and reports on each task to maximize their total utility, which is the difference between the reward received for their reports and the cost incurred in making evaluations. Formally, an agent i ’s strategy is a vector of D tuples $[(e_{ij}, f_{ij})]$, specifying her effort level e_{ij} as well as the function f_{ij} she uses to *map* her actual evaluation \hat{X}_{ij} into her report X_{ij} for each of her tasks. Note that since an agent’s proficiency on a task p_{ij} is a function of her strategically chosen effort e_{ij} , the proficiency of agent i for task j is *endogenous* in our model.

For a single task, we use the notation $(1, X)$ to denote the choice of full effort $e_{ij} = 1$ and truthfully reporting one’s evaluation (*i.e.*, f_{ij} is the identity function $X_{ij} = \hat{X}_{ij}$), $(1, X^c)$ to denote full effort followed by inverting one’s evaluation, and $(0, r)$ to denote the choice of exerting no effort ($e_{ij} = 0$) and simply reporting the outcome of a random coin toss with probability r of returning H . We use $[(1, X)]$ to denote the strategy of using full effort and truthtelling on all of an agent’s tasks, and similarly $[(1, X^c)]$ and $[(0, r)]$ for the other strategies.

Mechanisms. A mechanism in this setting takes as input the set of all received reports X_{ij} and computes a reward for each agent based on her reports, as well as possibly the reports of other agents.

Note that the mechanism has no access³ to the underlying true qualities \bar{X}_j for any task, and so cannot use the \bar{X}_j to determine agents’ rewards. A set of effort levels and reporting functions $[(e_{ij}, f_{ij})]$ is a full-information Nash equilibrium of a mechanism if no agent i can strictly improve her expected utility by choosing either a different effort level \hat{e}_{ij} , or a different function \hat{f}_{ij} to map her evaluation \hat{X}_{ij} into her report X_{ij} . Here, the expectation is over the randomness in agents’ noisy evaluation of the underlying ground truth, as well as any randomness in the mechanism.

We will be interested in designing mechanisms for which it is (i) an equilibrium for all agents to put in full effort and report their evaluations truthfully on all tasks, *i.e.*, use strategies $[(1, X)]$, and (ii) for which $[(1, X)]$ is the maximum utility (if not unique) equilibrium. We emphasize here that we do not address the problem of how to optimally aggregate the T reports X_{ij} for task j into a final estimate of \bar{X}_j : this is an orthogonal problem requiring application-specific modeling; our only goal is to elicit the best possible judgements to aggregate, by ensuring that agents find it most profitable to put in maximum effort into their evaluations and then report these evaluations truthfully.

3. MECHANISM

The main idea behind our mechanism \mathcal{M} is following. Recall that a mechanism does not have access to the true qualities \bar{X}_j , and therefore must compute rewards for agents that do not rely on directly observing \bar{X}_j . Since the only source of information about \bar{X}_j comes from the reports X_{ij} , a natural solution is to reward based on some form of agreement between different agents reporting on j , similar to the peer-prediction setting [12]. However, an easy way for agents to achieve perfect agreement with no effort is to always report H (or L). With just one task, it is difficult for a mechanism to distinguish between the scenario where agents achieve agreement by making accurate, high-effort, evaluations of the same ground truth, and the low-effort scenario where agents achieve agreement by always reporting H , especially if $\mathcal{P}[H]$ is high. However, in our setting, we have the benefit of *multiple* tasks and ratings, which could potentially be used to distinguish between these two strategies and appropriately reward agents to incentivize high effort.

\mathcal{M} uses the presence of multiple ratings to subtract out a *static* term B_{ij} from the agreement score, chosen so that there is no benefit to making reports that are independent of \bar{X}_j —roughly speaking, \mathcal{M} rewards an agent i for her report on task j for agreeing with another ‘reference’ agent $r_j(i)$ ’s report on the same task, but *only beyond* what would be expected if i and $r_j(i)$ were randomly tossing coins with their respective empirical frequencies of heads.

Let d denote the number of other reports made by i and $r_j(i)$ that are used in the computation of this statistic term B_{ij} based on the observed frequency of heads for each pair (i, j) . We use \mathcal{M}_d to denote the version of \mathcal{M} which uses d other reports from each of i and $r_j(i)$ to compute B_{ij} . To completely specify \mathcal{M}_d , we also need to specify a reference rater $r_j(i)$ as well as this set of d (non-overlapping) tasks performed by i and $r_j(i)$, for which we use the following notation. (We require these d other tasks to be non-overlapping so that the reports for these tasks X_{ik} and $X_{r_j(i)l}$

³Crowdsourcing is used typically precisely in scenarios where the number of tasks is too large for the principal (or a set of trusted agents chosen by the principal) to carry out herself, so it is at best feasible to verify the ground truth for a tiny fraction of all tasks, which fraction turns out to be inadequate (a formal statement is omitted here) to incentivize effort using knowledge of the \bar{X}_j .

are independent⁴, which is necessary to achieve the incentive properties of \mathcal{M}_d .)

DEFINITION 1 ($S_{ij}, S_{r_j(i)j}$). Consider agent i and task $j \in J(i)$, and a reference rater $r_j(i)$. Given a value of d ($1 \leq d \leq D-1$), let S_{ij} and $S_{r_j(i)j}$ be sets of d non-overlapping tasks other than task j performed by i and $r_j(i)$ respectively, i.e.,

$$S_{ij} \subseteq J(i) \setminus j, \quad S_{r_j(i)j} \subseteq J(r_j(i)) \setminus j, \quad S_{ij} \cap S_{r_j(i)j} = \emptyset, \\ |S_{ij}| = |S_{r_j(i)j}| = d.$$

A mechanism \mathcal{M}_d is completely specified by reference raters $r_j(i)$ and the sets S_{ij} and $S_{r_j(i)j}$, and rewards agents as defined below. Note that \mathcal{M}_d only uses agents' reports X_{ij} to compute rewards and not their maximum proficiencies p_i , which therefore need not be known to the system.

DEFINITION 2 (MECHANISM \mathcal{M}_d). \mathcal{M}_d computes an agent i 's reward for her report $X_{ij} \in \{0, 1\}$ on task j , R_{ij} , by comparing against a 'reference rater' $r_j(i)$'s report $X_{r_j(i)j}$ for j , as follows:

$$R_{ij} = A_{ij} - B_{ij}, \quad \text{where} \quad (1)$$

$$A_{ij} = X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j}), \quad \text{and}$$

$$B_{ij} = \left(\frac{\sum_{k \in S_{ij}} X_{ik}}{d} \right) \left(\frac{\sum_{l \in S_{r_j(i)j}} X_{r_j(i)l}}{d} \right) \\ + \left(1 - \frac{\sum_{k \in S_{ij}} X_{ik}}{d} \right) \left(1 - \frac{\sum_{l \in S_{r_j(i)j}} X_{r_j(i)l}}{d} \right), \quad (2)$$

where the sets S_{ij} and $S_{r_j(i)j}$ in B_{ij} are as in Definition 1. The final reward to an agent i is βR_i , where $R_i = \sum_{j \in J(i)} R_{ij}$ and β is simply a non-negative scaling parameter that is chosen based on agents' costs of effort.

The first term, A_{ij} , in R_{ij} is an 'agreement' reward, and is 1 when i and $r_j(i)$ both agree on their report, i.e., when $X_{ij} = X_{r_j(i)j} = 1$ or when $X_{ij} = X_{r_j(i)j} = 0$. The second term B_{ij} is the 'statistic' term which, roughly speaking, deducts from the agreement reward whatever part of i and $r_j(i)$'s agreement on task j is to be 'expected anyway' given their reporting statistics, i.e., the relative frequencies with which they report H and L . This deduction is what gives \mathcal{M} its nice incentive properties— while \mathcal{M} rewards agents for agreement via A_{ij} , \mathcal{M} also penalizes for blind agreement that agents achieve without effort, by subtracting out the B_{ij} term corresponding to the expected frequency of agreement if i and $r_j(i)$ were randomly choosing reports corresponding to their estimated means.

For example, suppose all agents were to always report H . Then A_{ij} is always 1, but $B_{ij} = 1$ as well so that the net reward is 0; similarly if agents chose their reports according to a random coin toss, even one with the 'correct' bias $\mathcal{P}[H]$, the value of A_{ij} is exactly equal to B_{ij} since there is no correlation between the reports for a particular task, again leading to a reward of 0. The reward function R_{ij} is designed so that it *only* rewards agents when they put in effort into their evaluations, which leads to the desirable incentive properties of \mathcal{M}_d . (We note that there are other natural statistics which might incentivize agents away from low-effort reports— e.g., rewarding reports which collectively have an empirical mean close to $\mathcal{P}[H]$, or for variance. However, it turns out that appropriately balancing the agreement term (which is necessary to ensure agents cannot simply report according to a coin toss with

⁴We assume that co-raters' identities are kept unknown to agents, so there is no collusion between i and $r_j(i)$.

bias $\mathcal{P}[H]$) with a term penalizing blind agreement to simultaneously ensure that $[(1, X)]$ is an equilibrium *and* the most desirable equilibrium is hard to accomplish.)

There are two natural choices for the parameter d , i.e., how many reports of i and $r_j(i)$ to include for estimating the statistic term that we subtract from the agreement score in R_{ij} ⁵. (i) In \mathcal{M}_{D-1} , we set $d = D-1$ and include all reports of agents i and $r_j(i)$, except those on their common task j . Here, the non-overlap requirement for sets S_{ij} and $S_{r_j(i)j}$ says that an agent i and her reference rater $r_j(i)$ for task j have *only* that task j in common. (ii) In \mathcal{M}_1 , we set $d = 1$, i.e., subtract away the correlation between the report of i and $r_j(i)$ on exactly one other non-overlapping task. In \mathcal{M}_1 , the non-overlap condition only requires that for each agent-task pair, there is a reference agent $r_j(i)$ available who has rated one other task that is different from the remaining tasks rated by i , a condition that is much easier to satisfy than that in \mathcal{M}_{D-1} . In § 4, we will see that \mathcal{M}_1 will require that the choices of (j, j') , where $\{j'\} = S_{ij}$ is the task used in the statistic term of i 's reward for task j , are such that each task j' performed by i is used exactly once to determine R_{ij} for $j \neq j'$. Note that this is always feasible, for instance by using task $j+1$ in the statistic term for task j for $j = 1, \dots, D-1$ and task 1 for task D .

4. ANALYZING \mathcal{M}

In this section, we analyze equilibrium behavior in \mathcal{M}_d .

4.1 Preliminaries

Recall that proficiency is the probability of correctly evaluating the true quality. We use $p[H]$ (respectively $p[L]$) to denote the probability that an agent observes H (respectively L) when making evaluations with proficiency p , i.e., the probability that $\hat{X}_{ij} = H$ is $p[H] = p\mathcal{P}[H] + (1-p)\mathcal{P}[L]$. Similarly, $q[H], q[L]$ and $p_i[H], p_i[L]$ correspond to the probabilities of seeing H and L when making evaluations with proficiencies q and p_i respectively.

Matrix representation of strategies. We will frequently need to consider the space of all possible strategies an agent may use in the equilibrium analysis of \mathcal{M}_d . While the choice of effort level e_{ij} in an agent's strategy $[(e_{ij}, f_{ij})]$ is easily described— there are only two possible effort levels 1 and 0— the space of functions f_{ij} through which an agent can map her evaluation \hat{X}_{ij} into her report X_{ij} is much larger. For instance, an agent could choose f_{ij} corresponding to making an evaluation, performing a Bayesian update of her prior on X_j , and choosing the report with the higher posterior probability. We now discuss a way to represent strategies that will allow us to easily describe the set of all reporting functions f_{ij} .

An agent i 's evaluation \hat{X}_{ij} can also be written as a two-dimensional vector $o^{ij} \in \mathbb{R}^2$, where $o^{ij} = [1 \ 0]^T$ if i observes a H , and $o^{ij} = [0 \ 1]^T$ if i observes a L , where \mathbf{a}^T denotes the transpose of \mathbf{a} . For the purpose of analyzing \mathcal{M}_d , any choice of reporting function f_{ij} can then be described via a 2×2 matrix

$$M^{ij} = \begin{bmatrix} x & 1-y \\ 1-x & y \end{bmatrix},$$

where x is the probability with which i chooses to report H after observing H , i.e., $x = \Pr(X_{ij} = H | \hat{X}_{ij} = H)$, and similarly $y = \Pr(X_{ij} = L | \hat{X}_{ij} = L)$. Observe that the choice of effort e_{ij} affects *only* o^{ij} and its 'correctness', or correlation with the (vec-

⁵Understanding the effect of the parameter d in our mechanisms, which appears irrelevant to the mechanism's behavior when agents are risk-neutral, is an interesting open question.

tor representing the) actual quality \bar{X}_j , and the choice of reporting function f_{ij} only affects M^{ij} .

Any reporting matrix M^{ij} of the form above can be written as a convex combination of four matrices— one for each of the f_{ij} corresponding to (i) truthful reporting ($X_{ij} = \hat{X}_{ij}$) (ii) inverting ($X_{ij} = \hat{X}_{ij}^c$), and (iii, iv) always reporting H or L independent of one's evaluation ($X_{ij} = H$ and $X_{ij} = L$ respectively):

$$M_X = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, M_{X^c} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, M_H = \begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}, M_L = \begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}.$$

That is, $M^{ij} = \alpha_1 M_X + \alpha_2 M_{X^c} + \alpha_3 M_H + \alpha_4 M_L$, where $\alpha_1 = x - \alpha_3$, $\alpha_2 = 1 - y - \alpha_3$, and $\alpha_3 = x - y$ and $\alpha_4 = 0$ if $x \geq y$, and $\alpha_3 = 0$ and $\alpha_4 = y - x$ if $y > x$. It is easily verified that $\alpha_i \geq 0$, and $\sum \alpha_i = 1$, so that this is a convex combination. Since all possible reporting strategies f_{ij} can be described by appropriately choosing the values of $x \in [0, 1]$ and $y \in [0, 1]$ in M^{ij} , every reporting function f_{ij} can be written as a convex combination of these four matrices.

The agent's final report X_{ij} is then described by the vector $M^{ij} o^{ij} \in \mathbb{R}^2$, where the first entry is the probability that i reports H , i.e., $X_{ij} = H$, and the second entry is the probability that she reports $X_{ij} = L$. The expected reward of agent i for task j can therefore be written using the matrix-vector representation (where T denotes transpose and $\mathbf{1}$ is the vector of all ones) as

$$\begin{aligned} E[R_{ij}] &= E[(M^{r_j(i)j} o^{r_j(i)j})^T M^{ij} o^{ij}] \\ &\quad + (\mathbf{1} - M^{r_j(i)j} o^{r_j(i)j})^T (\mathbf{1} - M^{ij} o^{ij}) \\ &\quad - [(M^{r_j(i)j} E[o^{r_j(i)j}])^T M^{ij} E[o^{ij}]] \\ &\quad + (\mathbf{1} - M^{r_j(i)j} E[o^{r_j(i)j}])^T (\mathbf{1} - M^{ij} E[o^{ij}]), \end{aligned}$$

which is linear in M^{ij} . So the payoff from an arbitrary reporting function f_{ij} can be written as the corresponding linear combination of the payoffs from each of the 'basis' functions (corresponding to M_X, M_{X^c}, M_H and M_L) constituting $f_{ij} = M^{ij}$. We will use this to argue that it is adequate to consider deviations to each of the remaining basis reporting functions and show that they yield strictly lower reward to establish that $[(1, X)]$ is an equilibrium of \mathcal{M}_d .

Equivalent strategies. For the equilibrium analysis, we will use the following simple facts. (i) The strategy $(0, X)$ (i.e., using zero effort but truthfully reporting one's evaluation) is equivalent to the strategy $(0, r)$ with $r = 1/2$, i.e., to the strategy of putting in no effort, and randomly reporting H or L independent of the evaluation \hat{X}_{ij} with probability $1/2$ each. (ii) The strategy $(1, r)$ is equivalent to the strategy $(0, r)$, since the report X_{ij} in both cases is completely independent of the evaluation \hat{X}_{ij} and therefore of e_{ij} .

Cost of effort. While agents do incur a higher cost when using $e_{ij} = 1$ as compared to $e_{ij} = 0$, we will not need to explicitly deal with the cost in the equilibrium analysis— if the expected reward from using a strategy where $e_{ij} = 1$ is strictly greater than the reward from any strategy with $e_{ij} = 0$, the rewards R_{ij} can always be scaled appropriately using the factor β (in Definition 2) to ensure that the net utility (reward minus cost) is strictly greater as well.

We note that bounds on the factor β could be estimated empirically without knowledge of the priors by estimating the cost of effort c_{ij} from maximum proficiencies obtained from a pre-screening (§2), by conducting a series of trials with increasing rewards and then using individual rationality to estimate the cost of effort from observed proficiencies in these trials.

Individual rationality and non-negativity of payments. The expected payments made by our mechanism to each agent are always nonnegative in the full-effort truthful reporting equilibrium, i.e., when all agents use strategies $[(1, X)]$. To ensure that the payments are also non-negative for every instance (of the tasks and reports) and not only in expectation, note that it suffices to add 1 to the payments currently specified, since the penalty term B_{ij} in the reward R_{ij} is bounded above by 1. We also note that individual rationality can be achieved by using a value of β large enough to ensure that the net utility $\beta R_{ij} - c(1)$ remains non-negative for all values of $\mathcal{P}[H]$ — while the expected payment R_{ij} does go to zero as $\mathcal{P}[H]$ tends to 1 (i.e., in the limit of vanishing uncertainty as the underlying ground truth is more and more likely to always be H (or always be L)), as long as there is some bound $\epsilon > 0$ such that $\max\{\mathcal{P}[H], \mathcal{P}[L]\} \leq 1 - \epsilon$, a simple calculation can be used to determine a value $\beta^*(\epsilon)$ such that the resulting mechanism with $\beta = \beta^*$ leads to nonnegative utilities for all agents in the full-effort truth-telling Nash equilibrium of \mathcal{M} .

4.2 Equilibrium analysis

We now analyze the equilibria of \mathcal{M}_d . We begin with a technical lemma and a definition.

LEMMA 3. *Let $f_\alpha(p, q) = pq + (1-p)(1-q) - \alpha(p[H]q[H] + (1-p[H])(1-q[H]))$. If $\alpha \leq 1$, (i) $f_\alpha(p, q)$ is strictly increasing in p if $q > 1/2$, and strictly increasing in q if $p > 1/2$. (ii) $f_\alpha(p, q)$ is nonnegative if $p, q \geq 1/2$, and positive if $p, q > 1/2$. (iii) Denote $f(p, q) \triangleq f_1(p, q)$. Then, $f(p, q) = f(q, p) = f(1-p, 1-q)$. Also $f(p, 1-q) = f(1-p, q) = -f(p, q)$.*

PROOF. Recall that $p[H] = p\mathcal{P}[H] + (1-p)\mathcal{P}[L]$, and similarly for $q[H]$.

$$\begin{aligned} f_\alpha(p, q) &= p(2q-1) + (1-q) \\ &\quad - \alpha(p\mathcal{P}[H] + (1-p)\mathcal{P}[L])(2q[H]-1) - (1-q[H]) \\ &= p[(2q-1) - \alpha(\mathcal{P}[H] - \mathcal{P}[L])(2q[H]-1)] + K_{-p} \\ &= p(2q-1)(1 - \alpha(\mathcal{P}[H] - \mathcal{P}[L])^2) + K_{-p}, \end{aligned}$$

where K_{-p} is a term that does not depend on p , and we use $2q[H]-1 = (2q-1)(\mathcal{P}[H] - \mathcal{P}[L])$ in the last step.

Note that $\mathcal{P}[H] - \mathcal{P}[L] < \mathcal{P}[H] < 1$ if $\max\{\mathcal{P}[H], \mathcal{P}[L]\} < 1$, so that $1 - \alpha(\mathcal{P}[H] - \mathcal{P}[L])^2 > 0$ if $\alpha \leq 1$. Therefore, $f_\alpha(p, q)$ is linear in p with strictly positive coefficient when $q > 1/2$ and $\alpha \leq 1$. An identical argument can be used for q since $f_\alpha(p, q)$ can be written as a linear function of q exactly as for p :

$$f_\alpha(p, q) = q(2p-1)(1 - \alpha(\mathcal{P}[H] - \mathcal{P}[L])^2) + K_{-q}.$$

This proves the first claim.

For nonnegativity of $f_\alpha(p, q)$ on $p \in [1/2, 1]$, we simply argue that $f_1(p, q)$ is increasing in q when $p \in [1/2, 1]$, and 0 at $q = 1/2$. So for any $q > 1/2$, $f_1(p, q) \geq 0$ for any $p \in [1/2, 1]$. But $f_\alpha(p, q)$ is decreasing in α , so $f_\alpha(p, q)$ is nonnegative for any $\alpha \leq 1$ as well.

The final claims about $f(p, q)$ and $f(1-p, q)$ can be verified just by substituting the definitions of $p[H]$ and $q[H]$ and from symmetry in p and q . \square

DEFINITION 4 (T_{ij}, d_{ij}). *Let T_{ij} be the set of all tasks $j' \neq j$ such that $j \in S_{ij'}$, i.e., T_{ij} is the set of tasks j' for which i 's report on task j is used to compute the statistic term of i 's reward $R_{ij'}$ for task j' . We use $d_{ij} = |T_{ij}|$ to denote the number of such tasks j' .*

Our main equilibrium result states that under a mild set of conditions on the choice of reference raters $r_j(i)$ and sets T_{ij} , exerting full effort and reporting truthfully on all tasks is an equilibrium of

\mathcal{M}_d —even when agents have different maximum proficiencies and can choose a different strategy for each task (for instance, an agent could choose to shirk effort on some tasks and put in effort on the others).

The main idea behind this result is best explained for the special case where all agents have the same maximum proficiency $p_i = p$ and are restricted to using the same strategy for each task. Here, the expected payoff from using $[(1, X)]$ turns out to be $f(p, p)$ where f is as defined in Lemma 3, while the payoff from playing $[(0, r)]$ is exactly 0 (independent of other agents' strategies); the payoff from deviating to $[(1, X^c)]$ when other agents play $[(1, X)]$ is $-f(p, p)$. Since $f(p, p) > 0$ for $p > 1/2$ and increases with p , it is a best response for every agent to attain maximum proficiency and truthfully report her evaluation. Extending the argument when agents can have *both* different maximum proficiencies p_i and use different strategy choices for each task requires more care, and are what necessitate the conditions on the task assignment in Theorem 5.

THEOREM 5. *Suppose $p_i > 1/2$ for all i , and for each agent i , for each task $j \in J(i)$, (i) $d_{ij} = d$, and (ii) $E[p_{r_j(i)}] = E_{j_l \in T_{ij}}[p_{r_{j_l}(i)}] \triangleq \bar{p}_i$, where the expectation is over the randomness in the assignment of reference raters to tasks and the sets T_{ij} . Then, $[(1, X)]$ is an equilibrium of \mathcal{M}_d .*

The first condition in Theorem 5, $d_{ij} = d$, says that each task j performed by an agent i must contribute to computing the reward via the statistic term for exactly d other tasks in $J(i)$, where d is the number of reports used to compute the ‘empirical frequency’ of H reports by i in the statistic term. The second condition $E[p_{r_j(i)}] = E_{j_l \in T_{ij}}[p_{r_{j_l}(i)}]$ says that an agent i should expect the average proficiency of her reference rater $r_j(i)$ to be equal for all the tasks that she performs, *i.e.*, agent i should not be able to identify any particular task where her reference raters are, on average, worse than the reference raters for her other tasks (intuitively, this can lead to agent i shirking effort on this task being a profitable deviation)⁶. The first condition holds for each of the two specific mechanisms \mathcal{M}_1 and \mathcal{M}_{D-1} , and the second condition can be satisfied by a randomization of the agents before assignment⁷. We now prove the result.

PROOF. Consider agent i , and suppose all other agents use strategy $[(1, X)]$, *i.e.*, put in full effort with truthtelling on all their tasks. It will be enough to consider pure strategy deviations, and show that there is no beneficial deviation to $(1, X^c)$, or $(0, r)$ for any $r \in [0, 1]$ on any single task or subset of tasks.

Consider a particular assignment of reference raters $r_j(i)$ and the sets S_{ij} (and T_{ij}). If the other agents all report according to $[(1, X)]$, then the total expected reward to agent i from all her D tasks in this assignment is

⁶We note that the conditions on \mathcal{M}_d (which can be easily satisfied for $d = 1$ and $D - 1$) arise because we allow *both* differing abilities p_i and choosing a different strategy for each task— if either generalization is waived, $[(1, X)]$ can be shown to be an equilibrium of \mathcal{M}_d even without imposing these conditions.

⁷The details for this randomized assignment of tasks to agents so that S_{ij} and $S_{r_j(i)j}$ are non-overlapping for all i, j and satisfy the conditions in Theorem 5 is included in the full version of the paper.

$$\begin{aligned} E[R_i] &= \sum_{j=1}^D E[X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j})] - \\ &\left[\frac{\sum_{k \in S_{ij}} E[X_{ik}]}{d} p_{r_j(i)}[H] + (1 - \frac{\sum_{k \in S_{ij}} E[X_{ik}]}{d}) (1 - p_{r_j(i)}[H]) \right] \\ &= \sum_{j=1}^D E[X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j})] - \\ &\left[\frac{\sum_{k \in S_{ij}} E[X_{ik}]}{d} (2p_{r_j(i)}[H] - 1) + (1 - p_{r_j(i)}[H]) \right] \\ &= \sum_{j=1}^D \left[E[X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j})] - \right. \\ &\left. \sum_{j_l \in T_{ij}} \left(\frac{E[X_{ij_l}]}{d} (2p_{r_{j_l}(i)}[H] - 1) - (1 - p_{r_j(i)}[H]) \right) \right], \end{aligned}$$

where the expectation is over any randomness in the strategy of i as well as randomness in i and $r_j(i)$'s evaluations for each task j , and we rearrange to collect X_{ij} terms in the last step.

Now, agent i can receive different reference raters and task sets S_{ij} in different assignments. So to compute her expected reward, agent i will also take an expectation over the randomness in the assignment of reference raters to tasks and the sets S_{ij} , which appear in the summation above via T_{ij} .

Recall the condition that $E[p_{r_j(i)}] = E_{j_l \in T_{ij}}[p_{r_{j_l}(i)}] \triangleq \bar{p}_i$. Using this condition and taking the expectation over the randomness in the assignments of $r_j(i)$ and S_{ij} , the expected reward of i is

$$\begin{aligned} E[R_i] &= \sum_{j=1}^D \left[E[X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j})] - \right. \\ &\left. \sum_{j_l \in T_{ij}} \left(\frac{E[X_{ij_l}]}{d} (2\bar{p}_i[H] - 1) - (1 - \bar{p}_i[H]) \right) \right] \\ &= \sum_{j=1}^D \left[E[X_{ij}X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j})] - \right. \\ &\left. \frac{d_{ij}}{d} E[X_{ij}] (2\bar{p}_i[H] - 1) - (1 - \bar{p}_i[H]) \right], \end{aligned}$$

where $\bar{p}_i[H] = E[p_{r_j(i)}[H]] = E_{j_l \in T_{ij}}[p_{r_{j_l}(i)}[H]]$.

The expected reward to agent i , when she makes evaluations with proficiency q_j for task j and truthfully reports these evaluations ($X_{ij} = \hat{X}_{ij}$), is then

$$\begin{aligned} E[R_i] &= \sum_{j=1}^D \left[q_j \bar{p}_i + (1 - q_j)(1 - \bar{p}_i) - \frac{d_{ij}}{d} q_j [H] (2\bar{p}_i[H] - 1) - (1 - \bar{p}_i[H]) \right] \\ &= \sum_{j=1}^D \left[q_j \bar{p}_i + (1 - q_j)(1 - \bar{p}_i) - \right. \\ &\left. \frac{d_{ij}}{d} (q_j [H] \bar{p}_i[H] + (1 - q_j [H])(1 - \bar{p}_i[H])) - (1 - \frac{d_{ij}}{d})(1 - \bar{p}_i[H]) \right] \\ &= \sum_{j=1}^D \left[f_{\frac{d_{ij}}{d}}(q_j, \bar{p}_i) + (\frac{d_{ij}}{d} - 1)(1 - \bar{p}_i[H]) \right]. \end{aligned} \quad (3)$$

where the expectation is taken over randomness in all agents' evaluations, as well as over randomness in the choices of $r_j(i)$ and S_{ij} .

We can now show that choosing full effort and truthtelling on all tasks is a best response when all other agents use $[(1, X)]$ if $d_{ij} = d$. First, by Lemma 3, $f_{\frac{d_{ij}}{d}}(q_j, \bar{p}_i)$ is increasing in q_j provided $\frac{d_{ij}}{d} \leq 1$, so agent i should choose full effort to achieve her

maximum proficiency p_i on all tasks. Next, note that in terms of the expected reward, using proficiency q_j and reporting X^c is equivalent to using proficiency $1 - q_j$ and reporting X . So again by Lemma 3, deviating to X^c , i.e., $(1 - q_j)$, on any task is strictly dominated by X for $q_j > 1/2$ and $\bar{p}_i > 1/2$.

Finally, if agent i chooses f_{ij} as the function which reports the outcome of a random coin toss with probability r of H for any task j , the component of $E[R_i]$ contributed by the term corresponding to X_{ij} becomes

$$\begin{aligned} & E \left[X_{ij} X_{r_j(i)j} + (1 - X_{ij})(1 - X_{r_j(i)j}) \right] - \\ & \frac{d_{ij}}{d} \left(E[X_{ij}] p_{r_j(i)}[H] + (1 - E[X_{ij}]) (1 - p_{r_j(i)}[H]) \right) \\ & = r p_{r_j(i)}[H] + (1 - r)(1 - p_{r_j(i)}[H]) - \\ & (r p_{r_j(i)}[H] + (1 - r)(1 - p_{r_j(i)}[H])) \\ & = 0, \end{aligned}$$

which is strictly smaller than the reward from $f_{ij} = X$ in (3) if $q_j > 1/2$ and $\frac{d_{ij}}{d} \geq 1$, since $f(q_j, \bar{p}_i)$ is strictly positive when $q_j, \bar{p}_i > 1/2$ by Lemma 3.

Since we need $\frac{d_{ij}}{d} \leq 1$ to ensure that $(1, X^c)$ is not a profitable deviation, and $\frac{d_{ij}}{d} \geq 1$ to ensure that $(0, r)$ is not a profitable deviation, requiring $d_{ij} = d$ simultaneously satisfies both conditions. Therefore, if $d_{ij} = d$, deviating from $(1, X)$ on any task j leads to a strict decrease in reward to agent i . Since the total reward to agent i can be decomposed into the sum of D terms which each depend only on the report X_{ij} and therefore the strategy for the single task j , any deviation from $[(1, X)]$ for any single task or subset of tasks strictly decreases i 's expected reward.

Therefore, the rewards $R_i = \sum_{j \in J(i)} R_{ij}$ can always be scaled appropriately to ensure that $[(1, X)]$ is an equilibrium of \mathcal{M}_d . \square

Other equilibria. While $[(1, X)]$ is an equilibrium, \mathcal{M}_d can have other equilibria as well—for instance, the strategy $[(0, r)]$, where all agents report the outcome of a random coin toss with bias r on each task, is also an equilibrium of \mathcal{M}_d for all $r \in [0, 1]$, albeit with 0 reward to each agent. In fact, as we show in the next theorem, no equilibrium, symmetric or asymmetric, in pure or mixed strategies, can yield higher reward⁸ to agents than $[(1, X)]$, as long as agents ‘treat tasks equally’ (for example, while an agent may choose to shirk effort on one task and work on all others, each of her tasks is equally likely to be the one she shirks on). We will refer to this as tasks being ‘apriori equivalent’, so that agents cannot distinguish between tasks prior to putting in effort on them (or equivalently, the assignment of reference raters is such that agents will not find it beneficial (in terms of expected reward) to use a different strategy for a specific task). Note that this assumption is particularly reasonable in the context of applications where agents are recruited for a collection of similar tasks as in crowdsourced abuse/adult content identification, or in peer grading where each task is an anonymous student’s solution to the same problem.

THEOREM 6. *Suppose $p_i > 1/2$, and tasks are apriori equivalent. Then, the equilibrium where all agents choose $[(1, X)]$ yields maximum reward to each agent.*

⁸We note that another equilibrium which achieves the same maximum expected reward is $[(1, X^c)]$, where all agents put in full effort to make their evaluations, but then all invert their evaluations for their reports. However, $[(1, X^c)]$ is a rather unnatural, and risky, strategy, and one that is unlikely to arise in practice. Also, as we will see later, $[(1, X^c)]$ can also lead to lower rewards when there are some agents who always report truthfully.

PROOF. Consider a particular agent i and task j , and a single potential reference rater $r_j(i)$ for (i, j) . Recall from the preliminaries that agent i 's choice of f_{ij} can be described via a matrix $M = \alpha_1 M_X + \alpha_2 M_{X^c} + \alpha_3 M_H + \alpha_4 M_L$, and that we denote i 's evaluation via a vector o , where $o = [1 \ 0]^T$ if i observes H and $o = [0 \ 1]^T$ if i observes L . Similarly, let us describe $r_j(i)$'s choice of reporting function via the matrix M' with corresponding coefficients α'_i , and denote $r_j(i)$'s evaluation by o' .

Since tasks are apriori equivalent, each player i (hence $r_j(i)$ too) uses strategies such that $E[X_{ij}] = E[X_{ik}]$ for all $j, k \in J(i)$. Then, we can rewrite the expected reward for agent i on task j , when paired with reference rater $r_j(i)$, as

$$E[R_{ij}] = 2(E[X_{ij} X_{r_j(i)j}] - E[X_{ij}] E[X_{r_j(i)j}]).$$

Using the matrix-vector representation, substituting M, M' with their representations in terms of the basis matrices and expanding, and evaluating the matrix-matrix products, we have

$$\begin{aligned} X_{ij} X_{r_j(i)j} &= o'^T M'^T M o = o'^T R_M o, \quad \text{where} \\ R_M &= (\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2) I + \alpha_2 \alpha'_1 M_{X^c} + \alpha_1 \alpha'_2 M_{X^c}^T \\ &+ (\alpha_3 \alpha'_3 + \alpha_4 \alpha'_4) \mathbf{1} + (\alpha_3 \alpha'_1 + \alpha_4 \alpha'_2) M_H + (\alpha_1 \alpha'_3 + \alpha_2 \alpha'_4) M_H^T \\ &+ (\alpha_4 \alpha'_1 + \alpha_3 \alpha'_2) M_L + (\alpha_1 \alpha'_4 + \alpha_2 \alpha'_3) M_L^T, \end{aligned}$$

and $I, \mathbf{1}$ denote the identity matrix and the matrix of all ones in $R_i^{2 \times 2}$ respectively, and we use $M_X^T M_X = M_{X^c}^T M_{X^c} = I$, $M_H^T M_H = M_L^T M_L = \mathbf{1}$, $M_{X^c}^T M_H = M_L$, $M_{X^c}^T M_L = M_H$, and $M_H^T M_L = \mathbf{0}$. Similarly,

$$E[X_{ij}] E[X_{r_j(i)j}] = E[o'^T M'^T] E[M o] = E[o'^T] R_M E[o],$$

where R_M is as defined above.

Now, note that $M_H o = [o_1 + o_2 \ 0]^T = [1 \ 0]^T$ since $o_1 + o_2 = 1$ for any evaluation vector o by definition, so that $E[o'^T M_H o] = E[o'^T] E[M_H o]$, since $M_H o$ is a constant. The same is the case for each of the terms $E[o'^T \mathbf{1} o]$, $E[o'^T M_{X^c}^T o]$, $E[o'^T M_L^T o]$, $E[o'^T M_L o]$. Therefore, these terms cancel out when taking the difference $E[X_{ij} X_{r_j(i)j}] - E[X_{ij}] E[X_{r_j(i)j}]$ (corresponding to the reward from either agent choosing to report X_{ij} independent of her evaluation being 0). Also note that $E[o'^j] = [p[H] \ p[L]]^T$ if agent i makes evaluations with proficiency p . Suppose the agents use effort leading to proficiencies p and p' respectively. Then, we have

$$\begin{aligned} & E[X_{ij} X_{r_j(i)j}] - E[X_{ij}] E[X_{r_j(i)j}] \\ &= (\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2) (E[o'^T o] - E[o']^T E[o]) \\ &+ \alpha_2 \alpha'_1 (E[o'^T M_{X^c} o] - E[o']^T M_{X^c} E[o]) \\ &+ \alpha_1 \alpha'_2 (E[o'^T M_{X^c}^T o] - E[o']^T M_{X^c}^T E[o]) \\ &= (\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2) (E[o'_1 o_1 + o_2 o'_2] - E[o'_1] E[o_1] - E[o_2] E[o'_2]) \\ &+ (\alpha_2 \alpha'_1 + \alpha_1 \alpha'_2) (E[o'_1 o_2 + o_1 o'_2] - E[o'_1] E[o_2] - E[o_1] E[o'_2]) \\ &= (\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2) \left(pp' + (1 - p)(1 - p') - p[H] p'[H] - \right. \\ &\quad \left. (1 - p[H])(1 - p'[H]) \right) + (\alpha_2 \alpha'_1 + \alpha_1 \alpha'_2) \left(p(1 - p') + \right. \\ &\quad \left. (1 - p)p' - p[H](1 - p'[H]) - (1 - p[H])p'[H] \right). \end{aligned}$$

Now, note that multiplier of $(\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2)$ is precisely $f(p, p')$, which by Lemma 3 is nonnegative if $p, p' \geq 1/2$, and strictly positive if $p, p' > 1/2$. Also, for $p, p' \geq 1/2$, note that $p[H] \leq p$ and $p'[H] \leq p'$. Now, the function $g(x, y) = x(1 - y) + y(1 - x)$ is decreasing in both x and y for $x, y \in [\frac{1}{2}, 1]$ (taking derivatives), so the multiplier of $(\alpha_2 \alpha'_1 + \alpha_1 \alpha'_2)$ is non-positive, and negative if $p, p' > 1/2$.

So the maximum value that $E[X_{ij} X_{r_j(i)j}] - E[X_{ij}] E[X_{r_j(i)j}]$ can take for nonnegative coefficients with $\sum \alpha_i = \sum \alpha'_i = 1$, is

$f(p, p')$, which is obtained by setting $\alpha_3 = \alpha_4 = 0$, $\alpha'_3 = \alpha'_4 = 0$ (*i.e.*, with no weight on random independent reporting), and $a_1 = \alpha'_1 = 1$, $\alpha_2 = \alpha'_2 = 0$ (or viceversa): this is because the maximum value of term $(\alpha_1 \alpha'_1 + \alpha_2 \alpha'_2)$ when $\alpha_2 = 1 - \alpha_1$ and $\alpha'_2 = 1 - \alpha'_1$ is 1 and is achieved with these values, which also minimize the value of the term $(\alpha_2 \alpha'_1 + \alpha_1 \alpha'_2)$ with the non-positive multiplier, since $(\alpha_2 \alpha'_1 + \alpha_1 \alpha'_2) \geq 0$ and is equal to 0 for these values of α_i, α'_i . Also, since $f(p, p')$ increases with increasing p and p' , it is maximized when agents put in full effort and achieve their maximum proficiencies $p_i, p_{r_j(i)}$.

Therefore the expected reward for the single component of $E[R_{ij}]$ coming from a specific reference rater achieves its upper bound when both agents use $[(1, X)]$. The same argument applies for each reference rater, and therefore to the expected reward $E[R_{ij}]$, and establishes the claim. \square

We next investigate what kinds of Nash equilibria might exist where agents use low effort with any positive probability. Apriori, it is reasonable to expect that there would be mixed-strategy equilibria where agents randomize between working and shirking, *i.e.*, put in effort (choose $e_{ij} = 1$) sometimes and not (choose $e_{ij} = 0$) some other times. However, we next show that as long as tasks are apriori equivalent and agents only randomize between reporting truthfully and reporting the outcome of an independent random cointoss (*i.e.*, they do not invert evaluations), the *only* equilibrium in which any agent uses any support on $(0, r)$ is the one in which *all* agents *always* use $(0, r)$ on all their tasks. To show this, we start with the following useful lemma saying that an agent who uses a low-effort strategy any fraction of the time will always have a beneficial deviation as long as some reference agent plays $(1, X)$ with some positive probability. Roughly speaking, this is because as long as there is some probability that an agent's reference rater plays $(1, X)$ rather than $(0, r)$, the agent strictly benefits by always playing $(1, X)$ to maximize the probability of both agents playing $(1, X)$, which is the only time the agent obtains a positive reward.

LEMMA 7. *Suppose the probability of agent i using strategy $(1, X)$ is δ and strategy $(0, r_i)$ is $1 - \delta$ for each task $j \in J(i)$. Suppose i 's potential reference raters $r_j(i)$ use strategies $(1, X)$ and $(0, r_{r_j(i)})$ with probabilities $\epsilon_{r_j(i)}$ and $1 - \epsilon_{r_j(i)}$ respectively, for each task $j \in J(i)$. If $\epsilon_{r_j(i)} > 0$ for any reference rater with proficiency $p_{r_j(i)} > 1/2$, then agent i has a (strict) profitable deviation to $\delta' = 1$, *i.e.*, to always using strategy $(1, X)$, for all values of $r_i \in [0, 1]$.*

PROOF. Consider a particular task j , and let $k = 1, \dots, K$ be the potential reference rater for (i, j) . Let a_k denote the probability that k is the reference rater for agent i for task j . By linearity of expectation, i 's expected reward for j can be written as

$$\begin{aligned} E[R_{ij}] &= \sum_{k=1}^K a_k \left[\delta \epsilon_k (p_i p_k + (1 - p_i)(1 - p_k)) \right. \\ &\quad - (p_i[H] p_k [H] + (1 - p_i[H])(1 - p_k)) \\ &\quad + (1 - \delta) \epsilon_k (r_i p_k [H] + (1 - r_i)(1 - p_k [H])) \\ &\quad - (r_i p_k [H] + (1 - r_i)(1 - p_k)) \\ &\quad + \delta (1 - \epsilon_k) (p_i [H] r_k + (1 - p_i [H])(1 - r_k)) \\ &\quad - (p_i [H] r_k + (1 - p_i [H])(1 - r_k)) \\ &\quad \left. + (1 - \delta) (1 - \epsilon_k) (r_i r_k + (1 - r_i)(1 - r_k)) - (r_i r_k + (1 - r_i)(1 - r_k)) \right] \\ &= \delta \sum_k a_k \epsilon_k (p_i p_k + (1 - p_i)(1 - p_k) \\ &\quad - (p_i [H] p_k [H] + (1 - p_i [H])(1 - p_k [H]))) \\ &= \delta \sum_k a_k \epsilon_k f_1(p_i, p_k). \end{aligned}$$

Now, $E[R_{ij}]$ is linear in δ , and by Lemma 3, the coefficient of δ is nonnegative for all ϵ_k and $p_k \geq 1/2$, and strictly greater than 0 if $\epsilon_k > 0$ for some k with $p_k > 1/2$. Therefore, i can strictly increase her expected reward $E[R_{ij}]$ by increasing δ for any $\delta < 1$, as long as there is some reference agent k with $\epsilon_k > 0$ and $p_k > 1/2$.

The same argument holds for each task $j \in J(i)$, and therefore to strictly improve i 's total reward $E[R_i]$, we only need one reference rater across all tasks to satisfy $\epsilon_k > 0$ and $p_k > 1/2$ to obtain a strictly beneficial deviation (recall that we assumed $p_i \geq 1/2$ for all i). \square

This lemma immediately allows us to show that the only low-effort equilibria of \mathcal{M} that we reasonably⁹ need to be concerned about is the pure-strategy equilibrium in which $e_{ij} = 0$ for all i, j . Note that different agents could use different r_i (or even r_{ij}) in such equilibria, but all agents will receive reward 0 in *all* such equilibria.

THEOREM 8. *Suppose every agent can be a reference rater with some non-zero probability for every other agent, and tasks are apriori equivalent. Then, the only equilibria (symmetric or asymmetric) in which agents mix between $(1, X)$ and any low-effort strategy $[(0, r_{ij})]$ with non-trivial support on $[(0, r_{ij})]$ are those where all agents always use low effort on all tasks.*

Eliminating low-effort equilibria. Our final result uses Lemma 7 to obtain a result about eliminating low-effort equilibria. Suppose there are some trusted agents (for example, an instructor or TA in the peer-grading context or workers with long histories of accurate evaluations or good performance in crowdsourcing platforms) who always report truthfully with proficiency $t > 1/2$. Let ϵ_t denote the minimum probability, over all agents i , that the reference rater for agent i is such a trusted agent (note that we can ensure $\epsilon_t > 0$ by having the trusted agent randomly choose each task with positive probability). Lemma 7 immediately gives us the following result for \mathcal{M}_a , arising from the fact that the reward from playing a random strategy $(0, r)$ is exactly 0—the presence of trusted agents with a non-zero probability, however small, is enough to eliminate low-effort equilibria altogether.

THEOREM 9. *Suppose $\epsilon_t > 0$. Then $[(0, r_{ij})]$ is not an equilibrium of \mathcal{M} for any $r_{ij} \in [0, 1]$.*

PROOF. Suppose all agents except the trusted agent use the strategy $(0, r_{ij})$, and ϵ_t is the probability that the trusted agent is the reference rater for any agent-task pair. Then, since agent i reports X_{ij} according to a random coin toss independent of the actual realization of j , the payoff from any reference rater, whether the trusted agent or another agent playing $(0, r)$ is 0. For notational simplicity, let $r = r_{ij}$, $r' = r_{r_j(i)j}$.

$$\begin{aligned} E[R_{ij}] &= \epsilon_t (rt[H] + (1 - r)(1 - t[H]) - (rt[H] + (1 - r)(1 - t[H]))) \\ &\quad + (1 - \epsilon_t) (rr' + (1 - r)(1 - r') - (rr' + (1 - r)(1 - r'))) \\ &= 0. \end{aligned}$$

By deviating to $(1, X)$, agent i can strictly improve her payoff as long as $\epsilon_t > 0$ and $t, p > 1/2$, since her expected reward from this deviation is

$$\begin{aligned} E[R_{ij}] &= \epsilon_t (pt + (1 - p)(1 - t) - (p[H]t[H] + (1 - p[H])(1 - t[H]))) \\ &\quad + (1 - \epsilon_t) (rr' + (1 - r)(1 - r') - (rr' + (1 - r)(1 - r'))) \\ &> 0, \end{aligned}$$

⁹(We say reasonably because of the technical possibility of equilibria where some agents mix over $(1, X^c)$ as well.)

since the coefficient of ϵ_t is positive for $t, p > 1/2$ by Lemma 3. Therefore, there is a strictly beneficial deviation to $(1, X)$, so there is a choice of multiplier for the reward such that the payoff to agent i , which is the difference between the reward and the cost of effort c , is strictly positive as well. So $(0, r_{ij})$ is not an equilibrium of \mathcal{M} when $\epsilon_t > 0$.

□

This result, while simple, is fairly strong: as long as some positive fraction of the population can be trusted to always report truthfully with proficiency greater than $1/2$, the only reasonable¹⁰ equilibrium of \mathcal{M} is the high-effort equilibrium $[(1, X)]$, no matter *how small* this fraction. In particular, note that \mathcal{M} does not need to assign a higher reward for agreement with a trusted agent to achieve this result, and therefore *does not need to know the identity* of the trusted agents. In contrast, the mechanism which rewards agents for agreement with a reference rater without subtracting out our statistic term must use a higher reward $w(\epsilon_t)$ for agreement with the trusted agents which increases as $\frac{1}{\epsilon_t}$ to eliminate low-effort equilibria¹¹—this, in addition to being undesirably large, also requires identification of trusted agents.

5. DISCUSSION

In this paper, we introduced the problem of information elicitation when agents' proficiencies are endogenously determined as a function of their effort, and presented a simple mechanism which uses the presence of multiple tasks to identify and penalize low-effort agreement to incentivize effort when tasks have binary types. Our mechanism has the property that maximum effort followed by truthful reporting is the Nash equilibrium with maximum payoff to all agents, including mixed strategy equilibria. In addition to handling endogenous agent proficiencies, to the best of our knowledge this is the first mechanism for information elicitation with this 'best Nash equilibrium' property over all pure and over mixed strategy equilibria that requires agents to only report their own evaluations (*i.e.*, without requiring 'prediction' reports of their beliefs about other agents' reports), and does not impose any requirement on a diverging number of agent reports per task to achieve its incentive properties. Our mechanism provides a starting point for designing information elicitation mechanisms for several crowdsourcing settings where proficiency is an endogenous, effort-dependent choice, such as image labeling, tagging, and peer grading in online education.

We use the simplest possible model that captures the complexities arising from strategically determined agent proficiencies, leading to a number of immediate directions for further work. First, our underlying outcome space is binary (H or L)—modeling and extending the mechanism to allow a richer space of outcomes and feedback is one of the most immediate and challenging directions for further work. Also, our model of effort is binary, where agents either exert full effort and achieve maximum proficiency, or exert

¹⁰ Again, we say reasonable rather than unique because $(1, X^c)$ does remain an equilibrium of \mathcal{M} for all ϵ_t less than a threshold value—however, in addition to being an unnatural and risky strategy, this equilibrium yields strictly smaller payoffs than $[(1, X)]$ when $\epsilon_t > 0$. Note also that the introduction of such trusted agents does not introduce new equilibria, and that $[(1, X)]$ remains an equilibrium of \mathcal{M} .

¹¹ The same is the case for a mechanism based on rewarding for the 'right' variance, which does retain $[(1, X)]$ as a maximum reward equilibrium, but still requires identifying the trusted agents and rewarding extra for agreement with them.

no effort to achieve the baseline proficiency. While our results extend to a model where proficiency increases linearly with cost, a natural question is how they extend to more general models, for example, with convex costs. Finally, a very interesting direction is that of *heterogenous* tasks with task-specific priors and abilities. In our model, tasks are homogenous with the same prior $\mathcal{P}[H]$, and agents have the same cost and maximum proficiency for each task. If tasks differ in difficulty, and agents can observe the difficulty of a task prior to putting in effort, there are clear incentives to shirk on harder tasks while putting in effort for the easier ones. While tasks are indeed apriori homogenous (or can be partitioned to be so) in some crowdsourcing settings, there are other applications where some tasks are clearly harder than others; also, agents may have task-specific abilities. Designing mechanisms with strong incentive properties for this setting is a very promising and important direction for further work.

Acknowledgements. We thank David Evans, Patrick Hummel and David Stavens for helpful discussions and pointers to related literature, and the anonymous referees for comments and suggestions that helped improve the presentation of the paper.

6. REFERENCES

- [1] P. Bolton and M. Dewatripont. *Contract theory*. The MIT Press, 2005.
- [2] Y. Chen, D. Bacon, I. Kash, D. Parkes, M. Rao, and M. Sridharan. Predicting your own effort. In *Proceedings of the 11th International Conference on Autonomous and Multiagent Systems (AAMAS 2012)*, 2012.
- [3] S. Goel, D. M. Reeves, and D. M. Pennock. Collective revelation: A mechanism for self-verified, weighted, and truthful predictions. In *Proceedings of the 10th ACM conference on Electronic commerce (EC 2009)*, pages 265–274. ACM.
- [4] C. Harris. You're Hired! An Examination of Crowdsourcing Incentive Models in Human Resource Tasks. In M. Lease, V. Carvalho, and E. Yilmaz, editors, *Proceedings of the Workshop on Crowdsourcing for Search and Data Mining (CSDM) at the Fourth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 15–18, Hong Kong, China, February 2011.
- [5] L. John, G. Loewenstein, and D. Prelec. Measuring the prevalence of questionable research practices with incentives for truth telling. *Psychological science*, 23(5):524–532, 2012.
- [6] R. Jurca and B. Faltings. Enforcing truthful strategies in incentive compatible reputation mechanisms. In *Proceedings of the 1st International Workshop on Internet and Network Economics (WINE 2005)*, volume 3828, pages 268–277. ACM, Springer Verlag.
- [7] R. Jurca and B. Faltings. Incentives for expressing opinions in online polls. In *Proceedings of the 9th ACM conference on Electronic commerce (EC 2008)*, pages 119–128, New York, NY, USA. ACM.
- [8] R. Jurca and B. Faltings. Minimum payments that reward honest reputation feedback. In *Proceedings of the 7th ACM conference on Electronic commerce (EC 2006)*, pages 190–199. ACM.
- [9] R. Jurca and B. Faltings. Robust incentive-compatible feedback payments. In *Trust, Reputation and Security: Theories and Practice*, volume 4452, pages 204–218. Springer-Verlag, 2007.

- [10] R. Jurca and B. Faltings. Mechanisms for making crowds truthful. *J. Artif. Int. Res.*, 34(1):209–253, Mar. 2009.
- [11] N. Lambert and Y. Shoham. Truthful surveys. *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE 2008)*, pages 154–165.
- [12] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting informative feedback: The peer-prediction method. *Management Science*, pages 1359–1373, 2005.
- [13] G. Paolacci, J. Chandler, and P. Ipeirotis. Running experiments on amazon mechanical turk. *Judgment and Decision Making*, 5(5):411–419, 2010.
- [14] A. Papakonstantinou, A. Rogers, E. Gerding, and N. Jennings. Mechanism design for the truthful elicitation of costly probabilistic estimates in distributed information systems. *Artificial Intelligence*, 175(2):648–672, 2011.
- [15] D. Prelec. A Bayesian Truth Serum for subjective data. *Science*, 306(5695):462–466, 2004.
- [16] D. Prelec and H. Seung. An algorithm that finds truth even if most people are wrong. *Working Paper*, 2007.
- [17] J. Witkowski and D. Parkes. A robust Bayesian Truth Serum for small populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI 2012)*.
- [18] J. Witkowski and D. C. Parkes. Peer prediction with private beliefs. In *Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC 2011)*.
- [19] J. Witkowski and D. C. Parkes. Peer Prediction without a Common Prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC 2012)*. ACM.