

Learning the Prior in Minimal Peer Prediction

JENS WITKOWSKI, Albert-Ludwigs-Universität Freiburg
DAVID C. PARKES, Harvard University

Many crowdsourcing applications rely on the truthful elicitation of information from workers; e.g., voting on the quality of an image label, or whether a website is inappropriate for an advertiser. Peer prediction provides a theoretical mechanism for eliciting truthful reports. However, its application depends on knowledge of a full probabilistic model: both a distribution on votes, and a posterior for each possible single vote received. In earlier work, Witkowski and Parkes [2012b], relax this requirement at the cost of “non-minimality,” i.e., users would need to both vote and report a belief about the vote of others. Other methods insist on minimality but still require knowledge of the distribution on votes, i.e., the signal prior but not the posterior [Jurca and Faltings 2008, 2011; Witkowski and Parkes 2012a]. In this paper, we develop the theoretical foundation for learning the signal prior in combination with these minimal peer-prediction methods. To score an agent, our mechanism uses the empirical frequency of reported signals against which to “shadow” [Witkowski and Parkes 2012a], delaying payments until the empirical frequency is accurate enough. We provide a bound on the number of samples required for the resulting mechanism to provide strict incentives for truthful reporting.

1. INTRODUCTION

Our focus in this paper is on a setting in which there are multiple tasks sampled from a stationary distribution. For example, these could be tasks in which workers are asked to vote on the appropriateness of content on a web page for advertisers. We refer to a particular task or item as a *world* and, to keep the presentation simple, we associate each agent with a single world. Each world is assumed to have a state with different states generating different observations by a worker (agent). Agents are asked to report their observation and we will score the report of an agent in a given world on the basis of (a) the empirical frequency of the reports of agents in other worlds, and (b) the report of other agents in the current world.

Our main result is a truthful mechanism that allows agents to have subjective beliefs about the prior probability of a particular world state and the conditional probability of an observation given a world state. All that is assumed to be common knowledge is a lower bound on the extent of the belief change from signal prior to signal posterior for any given agent. Moreover, we insist on *minimality*, i.e. agents are only asked to report their observation. Our mechanism is the first truthful mechanism that achieves both minimality and subjective prior beliefs.

The remainder of the paper is organized as follows: In Section 2, we introduce the model. After reviewing proper scoring rules and the classical peer prediction method in Section 3, we present *peer shadowing* in Section 4, a minimal peer prediction mechanism which requires only knowledge of a common signal prior but where the possible signal posteriors are allowed to be subjective. In Section 5, we then introduce *empirical peer shadowing* which allows agents to have subjective beliefs about both their prior and their possible posteriors, and where the signal prior required for peer shadowing is learned from agent reports in other worlds. In Section 5.2, we show how to use a form of Hoeffding’s inequality to derive upper bounds on the required number of samples given a lower bound on the extent of the belief change from signal prior to signal posterior. We conclude with remarks on ongoing work and directions for future work in Section 6.

Related Work

In addition to the original peer prediction method [Miller et al. 2005] that we will introduce in Section 3.2, there is other related work:

Jurca and Faltings [2007] apply techniques from robust optimization to the peer prediction method to make it robust against small variations in the commonly-held prior. Their work differs from ours in that we allow subjective priors to differ arbitrarily between agents, and in that the mechanism does not need to have a prior itself.

Prelec [2004] develops the “Bayesian truth serum” (BTS) for a setting where the prior need not be known to the mechanism, but must still be common to the agents. In addition to requiring a common prior, BTS is not minimal because it requires the agents to report both their observation and their belief about the observations of other agents.

For a setting with binary signals, Witkowski and Parkes [2012a] provide a Robust Bayesian Truth Serum (RBTS). As in Prelec’s mechanism, RBTS requires a common prior to agents but does not insist on the mechanism knowing the prior. Unlike Prelec’s mechanism, RBTS achieves strict incentive compatibility for every number of agents $n \geq 3$. The mechanism is based on the observation that a particularity of the quadratic scoring rule can be used to truthfully elicit signals even if the mechanism does not know an agent’s prior. This is the idea of identifying a “shadow” belief report by perturbing some other belief according to the agent’s signal report. In Section 4, we present “peer shadowing”, an application of this “shadowing method” to a peer prediction setting with semi-subjective priors (where the signal prior is common but the signal posteriors are not). In our main result of Section 5, we then use peer shadowing with a signal prior that is learned from reports of agents in other worlds.

Jurca and Faltings [2008; 2011] suggest a mechanism for *on-line* polls which is situated in the same common-prior setting as BTS and RBTS. Their mechanism is minimal, i.e. requires only a signal report, but it is not incentive compatible. Instead, the mechanism publishes the empirical frequencies of reports, and the authors show that these converge in equilibrium towards the true distribution of signals in the population. Their mechanism can also be regarded as a form of peer shadowing (Section 4) if the mechanism is assumed to know a common signal prior. Like our peer shadowing mechanism, their mechanism then allows for semi-subjective priors. However, peer shadowing is more robust than their “one-over-prior” mechanism which is crucial when coupled with learning: payments remain bounded for all learned signal priors and the mechanism is well-defined for learned signal priors of 0.

Most closely related to this paper is our earlier work [2012b], which is situated in the same setting without a common prior. We obtain incentive compatibility for every number of agents $n \geq 2$ (in the same world) and binary information. However, the mechanism in this earlier work is not minimal but requires the agents to report a belief report before making an observation. A crucial requirement in that work is “temporal separation,” i.e., the ability to elicit relevant information from an agent both before and after she makes her observation. The advantage of the combination of minimality and subjective priors is that agents are *allowed to hold complex subjective beliefs* potentially diverging from other agents’ beliefs, but *do not have to deliberate* about them since they only need to report their observation.

2. MODEL

Each item or task is referred to as a world. To keep the presentation simple we assume that each agent participates in a single world. The modification to allow each agent to participate in multiple worlds just requires care to only use worlds the agent does not participate in when calculating the empirical frequency with which to score that agent. Note that we cannot use any reports from worlds the agent participates in (not even from other agents) because, by participating, the agent learns something about the instantiated state of those worlds, and from the agent’s perspective signals would thus not be drawn according to the prior distribution anymore.

There are n different worlds, each of which contains at least two agents. Each agent is indexed such that agent i belongs to world i (modulo n). When interacting with its world, agent i observes a binary signal S_i , which is a random variable with values $\{0, 1\}$, that is sometimes represented $\{l, h\}$ and referred to as a “low” and a “high” signal, respectively. The signal represents an agent’s experience or opinion, and different world states induce different distributions on signals. The objective in peer prediction is to elicit an agent’s signal in an incentive compatible way.

Each agent i has a *subjective prior* in regard to the state of the world, and the signal it will receive conditioned on different world states. That is, every agent i has subjective beliefs in regard to a prior $\Pr_i(T = t)$ on the world state, in regard to the conditional probability $\Pr_i(S = h \mid T = t)$ for how signals are generated for each possible state t , and in regard to the number of possible states, denoted m_i . The mechanism does not need to know these priors and, moreover, the prior can vary from agent to agent. Note that an agent’s subjective prior is the agent’s belief about her own world before observing a signal and all other worlds for which she has not observed a signal. Collectively, we refer to an agent’s subjective beliefs as the agent’s *belief type*, denoted $\theta_i \in \Theta$ for some abstract set Θ . We insist that all belief types are admissible:

Definition 2.1 (Admissible belief). An agent’s belief type θ_i is admissible if the subjective prior satisfies the following properties:

- There are two or more possible states; i.e., $m_i \geq 2$
- Every state has positive probability, so that $\Pr_i(T = t) > 0$ for all $t \in \{1, \dots, m_i\}$.
- States are distinct, such that $\Pr_i(S = h \mid T = t) \neq \Pr_i(S = h \mid T = t')$ for any two $t \neq t'$.
- The signal beliefs conditional on state are fully mixed, with $0 < \Pr_i(S = h \mid T = t) < 1$ for all t .

We adopt the convention that states are sorted; i.e., $\Pr_i(S = h \mid T = 1) < \dots < \Pr_i(S = h \mid T = m_i)$.

Admissibility of an agent’s belief type is a weak requirement. In particular, note that any belief type can be transformed into an admissible belief type as long as (a) all signal beliefs conditional on state are fully mixed for states with positive probability, and (b) the signal beliefs conditional on state are distinct for at least two states with positive probability. Any two states with the same signal probability can be merged into a new state, and states with zero probability can be dropped.

When an agent observes a signal, she updates her world state and signal beliefs according to her subjective prior. It is important to emphasize that an agent’s subjective prior reflects her belief about the behavior of the true world and since the true world affects every agent in the same way, agent i has a subjective belief which is symmetric with regard to other agents. That is, agent i does not distinguish between her belief about an agent j and her belief about some other agent k . We can thus adopt shorthand $p_i(s_j | s_i) = \Pr_i(S_j = s_j \mid S_i = s_i)$ for agent i ’s posterior signal belief that any other agent j in the same world receives signal s_j given agent i ’s signal s_i .

The posterior signal belief can be calculated as

$$p_i(s_j | s_i) = \Pr_i(S_j = s_j \mid S_i = s_i) = \sum_{t=1}^{m_i} \Pr_i(S_j = s_j \mid T = t) \Pr_i(T = t \mid S_i = s_i), \quad (1)$$

and applying Bayes’ rule to the second part of the summation in (1) yields

$$\Pr_i(T = t \mid S_i = s_i) = \frac{\Pr_i(S_i = s_i \mid T = t) \Pr_i(T = t)}{\Pr_i(S_i = s_i)}. \quad (2)$$

The denominator in (2) is the prior signal belief and can be computed as

$$\Pr_i(S_i = s_i) = \sum_{t=1}^{m_i} \Pr_i(S_i = s_i \mid T = t) \Pr_i(T = t). \quad (3)$$

Similar to the posterior beliefs, we denote the prior signal belief for a high signal by $p_i(h) = \Pr_i(S_i = h)$. Note that $p_i(h)$ is agent i 's belief about (a) an agent in her own world observing a high signal before she observes a signal, i.e. $p_i(h) = \Pr_i(S_i = h) = \Pr_i(S_j = h)$, and (b) the distribution of high signals over all worlds. The latter is reflected in (3), where the latter part of the right hand side is the prior state belief of any randomly-picked world and former part of the right hand side is the probability of a high signal given that world (for $s_i = h$). Note that after a signal observation, agent i 's belief about the signals observed by other agents in her own world changes whereas $p_i(h)$ remains agent i 's belief about the distribution of high signals over all other worlds.

We will need Lemma 2.2 from earlier work. It states that after observing a high signal, an agent will increase her belief about another agent in the same world observing a high signal (decreasing after observing a low signal).

LEMMA 2.2. [Witkowski and Parkes 2012a] *For all admissible priors it holds that $1 > p_i(h|h) > p_i(h) > p_i(h|l) > 0$.*

The main requirement we place on the knowledge of the mechanism and the agents is that it the designer has a lower bound $b > 0$ on the distance between the prior $p_i(h)$ and the posteriors following $S_i = l$ and $S_i = h$, respectively, i.e.

$$b \leq \min(p_i(h) - p_i(h|l), p_i(h|h) - p_i(h)). \quad (4)$$

This provides a lower bound on how much the belief of an agent changes through observing a signal. It can be chosen arbitrarily small as long as it is strictly positive. In Section 5, we analyze the trade-off between low belief change bound b and low number of required samples.

3. PRELIMINARIES

In this section we review proper scoring rules and the classical peer prediction method.

3.1. Proper Scoring Rules

Proper scoring rules can be used to incentivize a rational agent to truthfully report her private belief about the likelihood of a future event.

Definition 3.1 (Binary Scoring Rule). Given possible outcomes $\Omega = \{0, 1\}$, and a report $y \in [0, 1]$ in regard to the probability of outcome $\omega = 1$, a binary scoring rule $R(y, \omega) : [0, 1] \times \{0, 1\} \rightarrow \mathbb{R} \cup \{-\infty, +\infty\}$ assigns a score based on report y and the outcome ω that occurs.

First, the agent is asked for her belief report $y \in [0, 1]$. Second, an event $\omega \in \{0, 1\}$ materializes (observed by the mechanism) and, third, the agent receives payment $R(y, \omega)$.

Definition 3.2 (Strictly Proper Scoring Rule). A binary scoring rule is *proper* if it leads to an agent maximizing her expected score by truthfully reporting her belief $p \in [0, 1]$ and *strictly proper* if the truthful report is the only report that maximizes the agent's expected score.

An example of a strictly proper scoring rule is the quadratic scoring rule.

Definition 3.3 (Quadratic Scoring Rule). The binary quadratic scoring rule R_q , normalized to give scores between 0 and 1, is given by:

$$\begin{aligned} R_q(y, \omega = 1) &= 2y - y^2 \\ R_q(y, \omega = 0) &= 1 - y^2. \end{aligned} \tag{5}$$

PROPOSITION 3.4. [e.g., Selten 1998] The binary quadratic scoring rule R_q is strictly proper.

We state the proof of the following lemma from Witkowski and Parkes [2012a] to build intuition for the analysis in the present paper.

LEMMA 3.5 (MINIMIZE DISTANCE). [Friedman 1983] Let $p \in [0, 1]$ be an agent's true belief about a binary future event. If the center scores the agent's belief report according to the quadratic scoring rule R_q but restricts the set of allowed reports to $Y \subseteq \mathbb{R}$, a rational agent will report a $y \in Y$ with minimal $(y - p)^2$ and thus minimal absolute difference $|y - p|$. The quadratic scoring rule is thus said to have quadratic loss.

PROOF. First observe that the quadratic scoring rule's two equations are well-defined for any $y \in \mathbb{R}$, including values of y outside $[0, 1]$. The expected score of reporting y if p is the true belief is $E[y] = p(2y - y^2) + (1 - p)(1 - y^2)$. Let's subtract this from the expected score given that the agent can submit a truthful report: $E[p] - E[y] = p(2p - p^2) + (1 - p)(1 - p^2) - p(2y - y^2) - (1 - p)(1 - y^2) = (p - y)^2$. Maximizing $E[y]$ is equivalent to minimizing $E[p] - E[y]$, and so we see that a rational agent will seek to minimize $(p - y)^2$ and thus minimize absolute difference $|p - y|$. \square

Friedman's property of effective scoring rules given the Euclidean metric is not satisfied by all scoring rules, and stated here for the quadratic scoring rule. Note that the lemma does not require the set Y to be contained in $[0, 1]$, and it thus holds for values y outside this range.

3.2. Classical Peer Prediction

The classical peer prediction method is defined for a *common prior*, shared by all agents and also known to the mechanism. In particular, it provides both a signal prior and a posterior for every possible signal. For this, we denote by $p(h|s_i) = \Pr(S_j = h \mid S_i = s_i)$ the signal posterior for a generic agent i that another generic agent j in the same world received a high signal given that agent i received signal s_i . While we present the binary version of the peer prediction method, it extends to an arbitrary number of signals.

The *classical peer prediction method* is defined as:

- (1) Each agent i is asked for her signal report $x_i \in \{0, 1\}$.
- (2) For each agent i , choose another agent j from the same world and pay agent i :

$$u_i = R(p(h|x_i), x_j), \tag{6}$$

where R is an arbitrary proper scoring rule and $x_j \in \{0, 1\}$ the signal report by agent j .

The mechanism knows the prior and thus can calculate $p(h|x_i)$ and the score for an agent.

THEOREM 3.6. [Miller et al. 2005] The classical peer prediction method is strictly Bayes-Nash incentive compatible for any strictly proper scoring rule R and any admissible common prior.




Fig. 1. Illustration of the peer shadowing mechanism with $p_i(h) \in (p_i(h|l), p_i(h|h))$. Note that $p_i(h|l)$ is closer to $y'_i = p_i(h) - d$ than to $y'_i = p_i(h) + d$, and that $p_i(h|h)$ is closer to $y'_i = p_i(h) + d$ than to $y'_i = p_i(h) - d$.

Example: Two agents each observe a signal in the same world, e.g. both agents inspect the same website to determine whether it contains inappropriate content. The common prior has two possible world states with $\Pr(T = 2) = 0.7$, as well as conditional signal probabilities $\Pr(S = h | T = 2) = 0.8$ and $\Pr(S = h | T = 1) = 0.1$. The prior probability that agent i will receive a high signal is therefore $\Pr(S_i = h) = \Pr(S_i = h | T = 2)\Pr(T = 2) + \Pr(S_i = h | T = 1)\Pr(T = 1) = 0.59$.

For the mechanism to be strictly Bayes-Nash incentive compatible, agent i 's unique best response should be truthful reporting when agent j is truthful. Agent i 's belief about the state following a high signal is $\Pr(T = 2 | S_i = h) = \Pr(S_i = h | T = 2)\Pr(T = 2) / \Pr(S_i = h) = 0.95$. The analogous update following a low signal leads to $\Pr(T = 2 | S_i = l) = 0.34$. Because of this belief update, agent i revises her belief that agent j received a high signal, with posterior signal beliefs $p(h|h) = \Pr(S_j = h | S_i = h) = 0.76$ and $p(h|l) = \Pr(S_j = h | s_i = l) = 0.34$. If agent i reports high, the center calculates signal posterior belief 0.76 and applies this, together with agent j 's signal report, to a strictly proper scoring rule. Agent i reporting her signal truthfully thus corresponds to her making a prediction about agent j 's signal report with her true belief about agent j 's signal. Assuming that agent j is truthful, agent i 's unique best response is thus to report truthfully.

4. PEER SHADOWING

We say that a peer prediction method is *minimal* if an agent's report is only a signal and not a belief. In this section, we adapt the shadowing method of Witkowski and Parkes [2012a] to provide a method for minimal peer prediction. Unlike classical peer prediction there is no need to determine a posterior for a given signal report. As in the "one-over-prior" mechanism of Jurca and Faltings [2008; 2011], peer shadowing requires only knowledge of a signal prior on the part of the designer. However, peer shadowing is more robust than the one-over-prior mechanism when coupled with learning: payments remain bounded for all learned signal priors $\hat{p}_i(h)$, and the mechanism is well-defined for learned signal priors of 0.

4.1. Mechanism

Let $S_j = s_j$ denote the signal observed by agent j in the same world as agent i . The *peer shadowing mechanism* is defined as follows (also see Figure 1):

- (1) Agent i receives a signal $S_i = s_i$ and, based on her belief type, she forms a posterior belief $p_i(h|s_i) \in \{p_i(h|l), p_i(h|h)\}$ about $S_j = h$.
- (2) The mechanism asks the agent for signal report $x_i \in \{0, 1\}$ and transforms it into a "shadow" posterior report

$$y'_i = \begin{cases} p(h) + d, & \text{if } x_i = 1 \\ p(h) - d, & \text{if } x_i = 0, \end{cases} \quad (7)$$

where $d > 0$ is a parameter of the mechanism.

- (3) The shadow posterior report y'_i , and agent j 's report are then applied to the quadratic scoring rule R_q to give agent i a score of

$$R_q(y'_i, x_j). \quad (8)$$

4.2. Analysis

Given that we allow for subjective priors, the solution concept we adopt is *ex post* subjective equilibrium [Witkowski and Parkes 2012b]. In this equilibrium concept, each agent i is best responding to the strategy of every other agent given common knowledge of rationality, common knowledge of admissible belief types, and knowledge of her own type (i.e., her own subjective prior). The equilibrium is *subjective* because it allows for each agent to have a distinct belief type, and *ex post* because it allows for strict uncertainty in regard to the types of other agents.

THEOREM 4.1 (TRUTHFULNESS). *The Peer Shadowing Mechanism is strictly ex post subjective incentive compatible given a common signal prior $p(h)$ known to the agents and the mechanism.*

PROOF. Given that agent j is truthful, we prove that agent i 's unique best response is to report truthfully. We establish this by reasoning about the distance between agent i 's signal posterior and shadow posterior. Without loss of generality, suppose agent i 's signal is $S_i = h$. There are two cases:

- $p(h) + d \leq p_i(h|h)$. But now $d > 0$, and so $p(h) - d < p(h) + d \leq p_i(h|h)$ and the result follows by Lemma 3.5.
- $p(h) + d > p_i(h|h)$. Since $p(h) < p_i(h|h)$ (Lemma 2.2) it holds that $(p(h) + d) - p_i(h|h) < p_i(h|h) - (p(h) - d)$ and the result follows by Lemma 3.5.

This completes the proof. \square

5. THE EMPIRICAL PEER SHADOWING MECHANISM

We are now ready to define an approach to adaptive peer shadowing that adopts an empirical frequency in place of the center's assumed knowledge of the signal prior. The empirical peer shadowing mechanism is minimal, i.e. it elicits reports consisting of only signals. Note that the mechanism withholds payments until every agent has reported her signal.

5.1. Mechanism

The *Empirical Peer Shadowing Mechanism* proceeds as follows:

- (1) First, every agent i reports her signal $x_i \in \{0, 1\}$ to the mechanism in private, i.e. without any other agent observing x_i .
- (2) Then, for each agent i :
 - (a) Let N be one agent from each world except agent i 's world. Compute the empirical frequency (empirical mean) $\hat{p}_i(h)$ of all $n - 1$ signal reports of agents in N :

$$\hat{p}_i(h) = \sum_{k \in N} \frac{x_k}{n - 1}.$$

- (b) Choose some $d > 0$ and compute the shadow posterior y'_i by “shadowing” from this empirical frequency:

$$y'_i = \begin{cases} \hat{p}_i(h) + d, & \text{if } x_i = 1 \\ \hat{p}_i(h) - d, & \text{if } x_i = 0, \end{cases} \quad (9)$$

- (c) Let j be another agent in the same world as agent i . Use shadowing to score agent i depending on how well y'_i predicts agent j 's signal:

$$u_i = R_q(y'_i, x_j)$$

5.2. Analysis

The main challenge in the analysis of the empirical peer-shadowing mechanism is to reason about the impact of whether or not the empirical frequency of high signals, $\hat{p}_i(h)$, lies “in between” the two possible posteriors for any finite number of samples. With some probability, we have $\hat{p}_i(h) \leq p_i(h|l)$ or $\hat{p}_i(h) \geq p_i(h|h)$, and in this case peer shadowing would not be truthful. We derive a lower bound on the expected benefit of being truthful given that $\hat{p}_i(h)$ lies within the interval and an upper bound on the expected benefit from a misreport when $\hat{p}_i(h)$ is outside the interval. Together with an upper bound on the probability that the empirical frequency is outside the interval, this provides a bound on the number of samples required for the empirical peer shadowing mechanism to have strict incentives for agents to be truthful.

We begin with two technical lemmas.

LEMMA 5.1. $p_i(h) - p_i(h|l) \leq p_i(h|h) - p_i(h) \Leftrightarrow p_i(h) \leq 0.5$.

PROOF. Note that

$$p_i(h|h) = 1 - p_i(l|h) = 1 - \frac{p_i(l)}{p_i(h)} p_i(h|l) = 1 - \frac{1 - p_i(h)}{p_i(h)} p_i(h|l) = 1 - \frac{p_i(h|l)}{p_i(h)} + p_i(h|l).$$

So we have

$$\begin{aligned} p_i(h) - p_i(h|l) \leq p_i(h|h) - p_i(h) &\Leftrightarrow p_i(h) - p_i(h|l) \leq 1 - \frac{p_i(h|l)}{p_i(h)} + p_i(h|l) - p_i(h) \\ \Leftrightarrow 2p_i(h) - 2p_i(h|l) + \frac{p_i(h|l)}{p_i(h)} - 1 \leq 0 &\Leftrightarrow 2p_i(h)^2 - 2p_i(h|l)p_i(h) + p_i(h|l) - p_i(h) \leq 0 \\ &\Leftrightarrow (2p_i(h) - 1) \underbrace{(p_i(h) - p_i(h|l))}_{>0} \leq 0 \Leftrightarrow p_i(h) \leq 0.5. \end{aligned}$$

This completes the proof. \square

LEMMA 5.2. *The smallest possible $p_i(h|h)$ given belief change bound b is:*

$$\underline{p}_i(h|h) = \begin{cases} 2\sqrt{b} - b, & \text{if } p_i(h) \leq 0.5 \\ 0.5 + b, & \text{if } p_i(h) \geq 0.5, \end{cases} \quad (10)$$

A lower bound for $p_i(h|h)$ given belief change bound b is $2\sqrt{b} - b$.

PROOF. We first prove the statement for $p_i(h) \geq 0.5$. From Lemma 5.1, we know that $b \leq p_i(h|h) - p_i(h)$ entails $b \leq p_i(h) - p_i(h|l)$, so that it is sufficient to minimize $p_i(h|h)$ subject to $b \leq p_i(h|h) - p_i(h)$. Then $p_i(h|h) = p_i(h) + b$ which is minimized for $p_i(h) = 0.5$, so that $\underline{p}_i(h|h) = 0.5 + b$ if $p_i(h) \geq 0.5$.

For the case where $p_i(h) \leq 0.5$, we can restrict the analysis to $b \leq p_i(h) - p_i(h|l)$ because we know from Lemma 5.1 that for $p_i(h) \leq 0.5$, $b \leq p_i(h) - p_i(h|l)$ entails $b \leq p_i(h|h) - p_i(h)$. From the proof of Lemma 5.1 we also know that

$$p_i(h|h) = 1 - \frac{p_i(h|l)}{p_i(h)} + p_i(h|l) = 1 - p_i(h|l) \left(\frac{1}{p_i(h)} - 1 \right) \quad (11)$$

Equivalent to minimizing $p_i(h|h)$ given $b \leq p_i(h) - p_i(h|l)$ is thus maximizing $p_i(h|l) \left(\frac{1}{p_i(h)} - 1 \right)$ given $b \leq p_i(h) - p_i(h|l)$. This is maximized for the largest possi-




Fig. 2. Illustration of Case 1 in the analysis of the empirical peer shadowing mechanism. Observe that $\epsilon < b$, so that $\hat{p}_i(h) = p_i(h) + \epsilon < p_i(h|h)$. (Depending on parameter $d > 0$, it may or may not hold that $\hat{p}_i(h) + d < p_i(h|h)$.)

ble $p_i(h|l)$ and the smallest possible $p_i(h)$, so that a necessary condition for a minimal $p_i(h|h)$ is $p_i(h|l) = p_i(h) - b$. Using this in (11) we obtain:

$$\underline{p}_i(h|h) = 1 - \left(\frac{p_i(h) - b}{p_i(h)} - (p_i(h) - b) \right) = \frac{b}{p_i(h)} + p_i(h) - b \quad (12)$$

Taking the derivative and setting to 0 one obtains:

$$\frac{\partial \underline{p}_i(h|h)(p_i(h))}{\partial p_i(h)} = 1 - \frac{b}{p_i(h)^2} = 0 \Leftrightarrow p_i(h) = \sqrt{b}$$

Inserting this back into (12), one obtains the minimal $p_i(h|h)$ for $p_i(h) \leq 0.5$:

$$\underline{p}_i(h|h) = \frac{b}{\sqrt{b}} + \sqrt{b} - b = 2\sqrt{b} - b. \quad (13)$$

Since $2\sqrt{b} - b \leq 0.5 + b$ for all $0 < b < 0.5$, this completes the proof. \square

In the proof of Theorem 5.4, we will use a form of Hoeffding's inequality in order to be able to make a statement about the number of samples we require without knowledge of $p_i(h)$. (For the simple steps showing how to get from the standard formulation to the formulation we use, see for example p.3 in Domke [2010].)

LEMMA 5.3. [Hoeffding 1963] Let $Z_1, \dots, Z_n \in [0, 1]$ be independent and identically distributed random variables. If

$$n \geq \frac{1}{2\epsilon^2} \ln \left(\frac{2}{\delta} \right),$$

for $\epsilon > 0$, $0 < \delta < 1$, then $\Pr(|\frac{1}{n} \sum_{i=1}^n Z_i - E[Z]| \leq \epsilon) \geq 1 - \delta$. That is, with probability at least $1 - \delta$, the difference between the empirical mean $\frac{1}{n} \sum_{i=1}^n Z_i$ and the expected value $E[Z]$ is at most ϵ .

THEOREM 5.4. The Empirical Peer-Shadowing Mechanism is strictly ex post subjective incentive compatible given belief change bound b and $n - 1$ samples with $n \geq \frac{1}{2\epsilon^2} \ln \left(\frac{2(1+2(b-\sqrt{b})-\epsilon)}{b-\epsilon} \right) + 2$ and $0 < \epsilon < b$.

PROOF. Given that all other agents are truthful, we show that agent i 's unique best response is to be truthful. To apply Hoeffding's inequality, we introduce some $\epsilon > 0$, $\epsilon < b$ and analyze two cases: the case where the empirical frequency is no more than ϵ away from the signal prior $p_i(h)$, i.e. $|\hat{p}_i(h) - p_i(h)| \leq \epsilon$, and the case where the empirical frequency is further away than ϵ , i.e. $|\hat{p}_i(h) - p_i(h)| > \epsilon$.

(Case 1): $|\hat{p}_i(h) - p_i(h)| \leq \epsilon$. (Also see Figure 2.)




Fig. 3. Illustration of Case 2 in the analysis of the empirical peer shadowing mechanism. Since $\hat{p}_i(h)$ is sampled using a finite number of samples and since it is not restricted to be within the ϵ bounds, it can happen that $\hat{p}_i(h) = 1$.

From $0 < \epsilon < b \leq \min(p_i(h) - p_i(h|l), p_i(h|h) - p_i(h))$ it follows that $p_i(h|l) < \hat{p}_i(h) < p_i(h|h)$, so that the shadowing method elicits signals truthfully. We proceed quantifying this positive expected benefit of reporting truthfully using the difference in expected loss given $S_i = h$ (case $S_i = l$ is analogous). Recall that the quadratic scoring rule has quadratic loss (Lemma 3.5):

$$\begin{aligned}
 & \Delta U_i(x_i = h | S_i = h) \\
 &= U_i(x_i = h | S_i = h) - U_i(x_i = l | S_i = h) \\
 &= - (p_i(h|h) - (\hat{p}_i(h) + d))^2 + (p_i(h|h) - (\hat{p}_i(h) - d))^2 \\
 &= \left((p_i(h|h) - (\hat{p}_i(h) - d)) + (p_i(h|h) - (\hat{p}_i(h) + d)) \right) \left((p_i(h|h) - (\hat{p}_i(h) - d)) - (p_i(h|h) - (\hat{p}_i(h) + d)) \right) \\
 &= (2p_i(h|h) - 2\hat{p}_i(h)) 2d = 4d (p_i(h|h) - \hat{p}_i(h))
 \end{aligned}$$

Using $p_i(h|h) \geq p_i(h) + b$ and $\hat{p}_i(h) \leq p_i(h) + \epsilon$, we derive lower bound

$$\Delta U_i(x_i = h | S_i = h) = 4d (p_i(h|h) - \hat{p}_i(h)) \geq 4d (p_i(h) + b - (p_i(h) + \epsilon)) = 4d (b - \epsilon)$$

on the gain in expected payment from reporting truthfully.

(Case 2): $|\hat{p}_i(h) - p_i(h)| > \epsilon$. (Also see Figure 3.)

In this case we provide an upper bound on the expected benefit from lying. Consider again without loss of generality that $S_i = h$:

$$\Delta U_i(x_i = l | S_i = h) = -\Delta U_i(x_i = h | S_i = h) = 4d (\hat{p}_i(h) - p_i(h|h)).$$

The maximal $\Delta U_i(x_i = l | S_i = h)$ is obtained for $\hat{p}_i(h)$ maximal and $p_i(h|h)$ minimal. Since $|\hat{p}_i(h) - p_i(h)| > \epsilon$, nothing prevents $\hat{p}_i(h) = 1$. From Lemma 5.2 we know that a lower bound of $p_i(h|h)$ given $p_i(h|h) - p_i(h) \geq b$ and $p_i(h) - p_i(h|l) \geq b$ is $2\sqrt{b} - b$, so that we can derive an upper bound for the expected benefit from lying by setting $\hat{p}_i(h) = 1$ and $p_i(h|h) = 2\sqrt{b} - b$, to obtain:

$$\Delta U_i(x_i = l | S_i = h) = 4d (\hat{p}_i(h) - p_i(h|h)) \leq 4d (1 - 2\sqrt{b} + b).$$

From Hoeffding's inequality, we know that Case 1 occurs with probability at least $1 - \delta$, so that the mechanism is truthful if

$$\begin{aligned}
(1 - \delta)4d(b - \epsilon) &> \delta 4d(1 - 2\sqrt{b} + b) \\
\Leftrightarrow (1 - \delta)(b - \epsilon) &> \delta(1 - 2\sqrt{b} + b) \\
\Leftrightarrow (b - \epsilon) &> \delta \underbrace{\left(1 + 2 \underbrace{(b - \sqrt{b})}_{\geq -0.25} - \underbrace{\epsilon}_{< b \leq 0.5}\right)}_{> 0} \\
\Leftrightarrow \delta &< \frac{b - \epsilon}{1 + 2(b - \sqrt{b}) - \epsilon}
\end{aligned}$$

To determine the number of worlds from which signals need to be sampled, the overall optimization problem becomes

$$\begin{aligned}
\text{min. } & n \\
\text{s.t. } & \epsilon < b \\
& \delta < \frac{b - \epsilon}{1 + 2(b - \sqrt{b}) - \epsilon} \\
& n - 1 \geq \frac{1}{2\epsilon^2} \ln\left(\frac{2}{\delta}\right).
\end{aligned}$$

The last line contains $n - 1$ instead of n because we compute $\hat{p}_i(h)$ using samples from $n - 1$ worlds. For any fixed ϵ , it is optimal to maximize δ , since this makes the right hand side of the final inequality as small as possible. Because of this, the problem can be restated as

$$\begin{aligned}
\text{min. } & n \\
\text{s.t. } & \epsilon < b \\
& \delta = \frac{b - \epsilon}{1 + 2(b - \sqrt{b}) - \epsilon} \\
& n - 1 > \frac{1}{2\epsilon^2} \ln\left(\frac{2}{\delta}\right),
\end{aligned}$$

where we have adopted equality for the second constraint and made the final inequality strict. Now, substituting for δ in the last inequality, we have:

$$\begin{aligned}
\text{min. } & \frac{1}{2\epsilon^2} \ln\left(\frac{2(1 + 2(b - \sqrt{b}) - \epsilon)}{b - \epsilon}\right) + 1 \\
\text{s.t. } & \epsilon < b
\end{aligned}$$

This completes the proof. \square

It is important to understand that the mechanism allows for subjective prior beliefs because it uses only "objective" signal reports which stem from the true world state to learn the prior. In particular, it does not elicit any beliefs from the agents. Since the signal reports used for learning the signal prior are not revealed to the agent, an agent forms a belief about this learned signal prior using her own belief type and it is therefore sufficient that the bounds we derive hold for any admissible belief type that satisfies belief change bound b .

Also note that for any given b , the minimal number of required samples can be computed numerically. For example, given bound $b = 0.05$, the optimal ϵ is $\epsilon = 0.046$, giving a corresponding requirement of $n - 1 = 1351$ samples. We believe sample numbers in this range are reasonable for applications such as eliciting votes on the quality of an

image label or whether a website is inappropriate for an advertiser. Note that these samples are from different items, so that we require that there are many images or websites and not that there are many votes on any particular image or website.

6. CONCLUSION

In this paper, we have presented an incentive compatible peer prediction mechanism which is the first to combine minimality with priors that are fully subjective and known only to the respective agent. This combination is compelling because it provides robustness against agents with non-standard (and possibly wrong) beliefs without imposing cognitive costs onto agents intending to report truthfully. In the analysis of the Empirical Peer Shadowing Mechanism, we derive an upper bound on the number of worlds (items) one needs to sample from. We believe that our mechanism could already be applied in applications such as crowdsourced image tagging, where requesters elicit information about many different items.

In addition to tightening the current analysis in regard to the number of samples required for strict incentives, and obtaining analytical bounds that are stated just in terms of b and not ϵ and b , there are two major future directions for this work. The first is to extend the mechanism to work in a setting with multiple signals. Second, we plan on designing a truthful mechanism for the orthogonal setting presented by opinion polls where the mechanism has access to many signal reports from the same world for learning information that can be used to incentivize truthful reporting but where the signal reports are all coming from this one world.

REFERENCES

- DOMKE, J. 2010. Statistical Machine Learning. <http://users.cecs.anu.edu.au/~jdomke/courses/sml2010/10theory.pdf>. Lecture Notes, Learning Theory (10).
- FRIEDMAN, D. 1983. Effective Scoring Rules for Probabilistic Forecasts. *Management Science* 29, 4, 447–454.
- HOEFFDING, W. 1963. Probability Inequalities for Sums of Bounded Random Variables. *Journal of the American Statistical Association* 58, 301, 13–30.
- JURCA, R. AND FALTINGS, B. 2007. Robust Incentive-Compatible Feedback Payments. In *Trust, Reputation and Security: Theories and Practice*. LNAI Series, vol. 4452. Springer-Verlag, 204–218.
- JURCA, R. AND FALTINGS, B. 2008. Incentives for Expressing Opinions in Online Polls. In *Proceedings of the 9th ACM Conference on Electronic Commerce (EC'08)*.
- JURCA, R. AND FALTINGS, B. 2011. Incentives for Answering Hypothetical Questions. In *Proceedings of the 1st Workshop on Social Computing and User Generated Content (SC'11)*.
- MILLER, N., RESNICK, P., AND ZECKHAUSER, R. 2005. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51, 9, 1359–1373.
- PRELEC, D. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306, 5695, 462–466.
- SELTEN, R. 1998. Axiomatic Characterization of the Quadratic Scoring Rule. *Experimental Economics* 1, 43–61.
- WITKOWSKI, J. AND PARKES, D. C. 2012a. A Robust Bayesian Truth Serum for Small Populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*.
- WITKOWSKI, J. AND PARKES, D. C. 2012b. Peer Prediction Without a Common Prior. In *Proceedings of the 13th ACM Conference on Electronic Commerce (EC'12)*.