

Peer Prediction without a Common Prior[†]

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

Reputation mechanisms at online opinion forums, such as Amazon Reviews, elicit ratings from users about their experience with different products. Crowdsourcing applications, such as image tagging on Amazon Mechanical Turk, elicit votes from users as to whether or not a job was duly completed. An important property in both settings is that the feedback received from users (agents) is truthful. The *peer prediction method* introduced by Miller et al. [2005] is a prominent theoretical mechanism for the truthful elicitation of reports. However, a significant obstacle to its application is that it critically depends on the assumption of a common prior amongst both the agents and the mechanism. In this paper, we develop a peer prediction mechanism for settings where the agents hold subjective and private beliefs about the state of the world and the likelihood of a positive signal given a particular state. Our *shadow peer prediction mechanism* exploits temporal structure in order to elicit two reports, a belief report and then a signal report, and it provides strict incentives for truthful reporting as long as the effect an agent's signal has on her posterior belief is bounded away from zero. Alternatively, this technical requirement on beliefs can be dispensed with by a modification in which the second report is a belief report rather than a signal report.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social and Behavioral Sciences—Economics; K.4 [Computers and Society]: Electronic Commerce

Additional Key Words and Phrases: Peer Prediction, Information Elicitation, Mechanism Design

1. INTRODUCTION

User-generated content is essential to the effective functioning of many e-commerce platforms, with prominent examples including the elicitation of feedback about products or services on sites such as Amazon Reviews or Expedia, and the elicitation of information or opinions from users (workers) on crowdsourcing platforms who are paid small rewards to do human computation tasks. While statistical estimation techniques [Raykar et al. 2010] can be adopted for the purpose of adjusting for biases or identifying users whose inputs are especially noisy, they are appropriate only in settings with repeated participation by the same user and when user inputs are informative in the first place. But what if providing accurate information is costly for users, or if users have external incentives for submitting false reports?

Consider a worker who provides feedback on the appropriateness of search results for given key words, who would rather shirk and avoid expending effort by randomly selecting a level of appropriateness. Alternatively, consider a worker in a crowdsourc-

[†]This paper is significantly extended from an earlier paper in the Workshop on Social Computing and User Generated Content (2011).

This work is supported in part by NSF Grant No. CCF-0915016 and a Microsoft Faculty Grant. Jens Witkowski is also grateful for financial support through Ph.D. fellowships from the Landesgraduiertenförderung Baden-Württemberg and the German Academic Exchange Service.

Author's addresses: J. Witkowski, Computer Science Department, Albert-Ludwigs-Universität Freiburg; D.C. Parkes, School of Engineering and Applied Sciences, Harvard University.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

EC'12, June 4–8, 2012, Valencia, Spain.

Copyright 2012 ACM 978-1-4503-1415-2/12/06...\$10.00.

ing context, such as voting on image tags or labeling websites that contain inappropriate content for an advertiser, who is paid per task and can thus improve his hourly rate by skipping over tasks without exerting due effort. Or consider a user in a restaurant rating context who may be concerned that sharing truthful, positive feedback could cause a trend towards the restaurant becoming more popular, and thus crowded.

The *peer prediction method* developed by Miller, Resnick and Zeckhauser [2005] (henceforth, MRZ) provides a theoretical approach to aligning incentives with reporting truthful feedback, or alternatively, with investing effort in order to provide informative feedback. Each user receives a payment from the center, with the amount of the payment depending on the user's own report and the report that was given by some other user. The information (or opinion) available to one user tells her something about the information (or opinion) that will be attained by another user. Take a digital camera bought from Amazon as an example: while different buyers may experience different quality, for example due to variations in the production process, all buyers receive the identical model. A buyer's belief that another buyer of the same camera experiences high quality is higher if she herself experienced high quality than if she experienced low quality.

The peer prediction method has fostered much research, including the work by Jurca and Faltings [2006; 2007; 2009] and Witkowski [2009; 2010]. To the best of our knowledge, however, it has not yet been applied in practice. The focus of the present paper is to address what we believe is the most significant difficulty when turning to application: the peer prediction method makes strong common knowledge assumptions. In particular, in the context of eliciting feedback on the quality of a product or service, it requires that every user shares the same prior belief about an item's inherent quality and about the way in which a user's experience (or "signal") is generated from quality. Moreover, it assumes that the mechanism knows this prior.

To support these assumptions, MRZ suggest that, in the context of feedback about a product or service, the rating history can be leveraged in order to allow the mechanism to estimate the prior. This, however, leaves open the question as to how the rating history itself was built: either people reported honestly without an incentive-compatible mechanism in place (and there is no design problem) or people reported dishonestly, in which case the mechanism cannot use these reports to learn the correct prior. Beyond this difficulty, there remains the concern that every user must share the same prior, something that seems unreasonable in practice.

Relaxing this common prior assumption is the main focus of our paper, and we introduce two mechanisms that provide strict incentives for truthful reports in equilibrium while allowing participants to adopt subjective and private priors. Both mechanisms ask a user for two reports: one before experiencing an item (or in crowdsourcing, before the user is presented with a unit of work), and one afterwards. The ability to enforce this *temporal separation* is critical but seems very reasonable in many applications. For example, a travel site could ask a user for her opinion at the time of booking and then again after her stay. Similarly, a paper reviewing system could ask a reviewer about her opinion on a particular paper after providing access to the paper's title and abstract, and then again after providing access to the entire paper.

We first introduce the *basic private-prior peer-prediction mechanism (BPP)*, that requires an agent to report two *belief* reports about the signal that another agent will receive, one before and one after receiving her own signal. BPP is strictly incentive compatible, and infers the agent's signal from the change in her belief reports. Building on this, we introduce and analyze the equilibrium of a mechanism in which an agent's belief report is followed by only a *signal* report. In this candidate mechanism, truthful reporting of the signal, but not of the belief, is an equilibrium. Computing an agent's optimal belief misreport, we then construct the strictly incentive compatible

shadow private-prior peer-prediction mechanism (SPP) via an application of the *revelation principle*, simulating this misreport on behalf of an agent. Moreover, we present a special case of SPP that has a very simple and intuitive form. The main technical innovation is to combine an agent’s belief report and her own signal report to calculate a “shadow” posterior belief which is scored according to its predictiveness of another agent’s signal report. The only knowledge SPP requires about the agents’ priors is that the effect an agent’s signal has on her posterior belief that another agent will receive a high signal is bounded away from zero. In particular, the mechanism requires knowledge of some $\epsilon > 0$ such that the difference between the posterior belief following a high signal and the posterior belief following a low signal, is bounded from below by ϵ . This technical requirement in regard to the minimal informativeness of signals given the agents’ beliefs is not required for the strict incentives of BPP, in which the second report is a belief report rather than a signal report.

In moving from a common knowledge setting to one with private and subjective priors, an important consideration is the amount of additional information over and above a signal report that must be elicited from a participant. Indeed, MRZ had suggested the possibility of incentive compatible peer prediction with subjective and private priors. In a brief treatment, they proposed an approach in which in addition to her own signal, a user also reports her prior belief on the true underlying world state (e.g., a product’s “quality”) and her belief on the probability of receiving each possible signal, conditioned on each possible state. In comparison, both BPP and SPP are considerably simpler with respect to the reporting costs, and thus likely more practical. In fact, our analysis suggests a trade-off between the robustness of incentive properties and the reporting requirements, given that SPP but not BPP requires the technical requirement on the minimal informativeness of signals.

A limitation of both BPP and SPP relative to the original peer prediction method is that its application is restricted to domains with *binary* signals. However, many interesting applications of peer prediction mechanisms are to settings with binary signals. For example, blogs and online forums allow users to vote whether a post was helpful or not. Similarly, social networking websites, such as Facebook and Google+, allow users to “like” or “+1” other users’ comments. Hotel booking websites, such as Expedia and Hotwire, ask customers after their stay whether they “would recommend this hotel to a friend,” and the decision whether to flag a crowdsourced task as spam or not, is binary, too. While we believe that it is an interesting theoretical question whether BPP and SPP can be extended to handle multiple signals, note that, even with their relatively low reporting costs, this would still place a high burden on users, who, when reporting their beliefs, would have to report full probability distributions instead of a single number.

Related Work

In addition to the original peer prediction method [Miller et al. 2005] that we will introduce in Section 3, 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 incentive compatible only for a large enough number of agents. For a setting with binary signals, Witkowski and Parkes [2012] 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 reference belief report according to the agent’s signal report. However, the design problem for the private and *subjective* prior setting of this paper is significantly more difficult than the common-prior setting of RBTS. In particular, whereas another agent provides the reference belief in RBTS, in the present context with subjective priors, an agent’s *own* belief report must provide the reference. This complicates the design problem significantly, because there is now a coupling between the belief report and the signal report.

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 requires only a signal report but it is not incentive compatible. Instead, the mechanism publishes the current frequencies of reports, and the authors show that these converge in equilibrium towards the true distribution of signals in the population.

In handling private and subjective priors, in place of the Bayes-Nash equilibrium analysis of peer prediction mechanisms, we analyze the *subjective equilibrium* of a mechanism. Informally, this requires that each agent best-responds to the strategy of every other agent given her own subjective prior, and given strict uncertainty about the priors of other agents. Formally, the prior of an agent is modeled via a *belief type* and we study the properties of BPP and SPP in an *ex post subjective equilibrium*, emphasizing that agents need no knowledge of other agents’ belief types. Various notions of subjective, self-confirming and conjectural equilibria appear in the game theory literature, although normally studied in repeated contexts [e. g., Kalai and Lehrer 1993; P. Battigali and Molinari 1992]. The most related concept to our solution concept is that of Rubinstein and Wolinsky [1984], who propose *rationalizable conjectural equilibria* in the context of a one-shot game, but without our notions of *ex post* robustness and strict uncertainty about the others’ belief types. In addition, they require player observations in regard to the play of others and consistency of beliefs relative to these observations. We model a single interaction between an agent and a peer prediction mechanism, and the agent is unable to make any observations about the actions of her “peer” agent (against which she is scored.)

2. THE SETTING

There is a group of $n \geq 2$ rational, risk-neutral and self-interested agents. A *world state* is determined by random variable T which can adopt values in $\{1, \dots, m\}$. When interacting with the world, each 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 signals 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, a prior can vary from agent to agent. 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. An agent’s belief type 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 $\Pr_i(S = h | T = t) \neq \Pr_i(S = h | T = t')$ for any two $t \neq t'$. We adopt the convention that states are sorted; i.e., $\Pr_i(S = h | T = 1) < \dots < \Pr_i(S = h | T = m_i)$.
- The signal beliefs conditional on state are fully mixed, with $0 < \Pr_i(S = h | T = t) < 1$ for all t .

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 state and signal beliefs according to her subjective prior. We adopt shorthand $p_i(s_i) = \Pr_i(S_j = h | S_i = s_i)$ for agent i ’s posterior signal belief that a second agent j receives a high signal given agent i ’s signal s_i . The posterior signal belief can be calculated as

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

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

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

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

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

Similar to the posterior beliefs, we denote the prior signal belief for a high signal by $p_i = \Pr_i(S_i = h)$.

An important property that we require of the domain is temporal separation:

Definition 2.2 (Temporal Separation). A domain satisfies temporal separation if a mechanism can restrict one report from an agent to being made before the agent receives her signal and a second report to being made after the agent receives her signal.

At this stage, we are not specific about what constitutes a “report” but we are heading towards a mechanism, in which a first report about an agent’s prior signal belief p_i is separated from the agent’s report about her received signal, where it is crucial that the report in regard to p_i must be made before the agent receives her signal.

3. THE PEER PREDICTION METHOD

In this section we briefly review proper scoring rules and the classical peer prediction method. 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 $R(y, \omega) \in \mathbb{R} \cup \pm \infty$ 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 binary logarithmic scoring rule R_l :

$$\begin{aligned} R_l(y, \omega = 1) &= \ln(y) \\ R_l(y, \omega = 0) &= \ln(1 - y), \end{aligned} \tag{4}$$

for natural logarithm \ln .

PROPOSITION 3.3. [e.g., Selten 1998] *The binary logarithmic scoring rule R_l is strictly proper.*

Note that if one applies a positive-affine transformation to a proper scoring rule, the rule is still proper. For a more detailed discussion of proper scoring rules in general, we refer to the article by Gneiting and Raftery [2007].

The classical peer prediction method is defined for a *common prior*, shared by all agents and also known to the mechanism. For this, we denote by $p(s_i) = \Pr(S_j = h \mid S_i = s_i)$ the signal posterior for a generic agent i , given that it has 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 $j = i + 1$ (modulo n) and pay agent i :

$$u_i = R(p(x_i), x_j), \tag{5}$$

where R is an arbitrary proper scoring rule and x_j the signal report by agent j .

Because the common prior is known to the mechanism, an agent's score $R(p(x_i), x_j)$ can be computed by the center for all $x_i, x_j \in \{0, 1\}$.

Consider a strategy $\sigma_i : \{0, 1\} \rightarrow \{0, 1\}$ for an agent, describing her report for every possible signal. Given a common prior, the equilibrium concept adopted for the analysis of peer prediction is that of a Bayes-Nash equilibrium:

Definition 3.4. Strategy profile $(\sigma_1, \dots, \sigma_n)$ is a *Bayes-Nash equilibrium (BNE)* of the peer prediction method with n agents if

$$E[R(p(\sigma_i(s_i)), \sigma_j(S_j)) \mid S_i = s_i] \geq E[R(p(s'_i), \sigma_j(S_j)) \mid S_i = s_i],$$

for all i , all s_i , all $s'_i \neq s_i$, and with $j = i + 1$ (modulo n).

Each agent maximizes her expected score by following strategy σ_i , given her own signal, her posterior signal belief, and given that the other agents play according to strategy profile σ_{-i} . A mechanism is *Bayes-Nash incentive compatible (BNIC)* if the *truthful strategy*, $\sigma_i(s_i) = s_i$ for each $s_i \in \{0, 1\}$, and for all agents i , is a BNE. A mechanism is *strictly BNIC* if the inequalities defining the BNE are strict rather than weak, so that agents strictly prefer to submit a truthful report.

THEOREM 3.5. [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.*

Definition 3.6 (Ex Post Individually Rational). A peer prediction mechanism is *ex post individually rational (ex post IR)* if no agent can ever incur a negative payment.

PROPOSITION 3.7. [Miller et al. 2005] *The classical peer prediction method can be made ex post IR for any proper scoring rule R .*

In the classical peer prediction method, the mechanism knows the prior and thus all possible posteriors and scores that can arise for a given scoring rule. It is thus possible to add a constant to every payment, such that, for example, the logarithmic rule plus this constant is non-negative for any possible combination of signal reports.

Example 1. (Classical Peer Prediction) Two agents each experience a product that can have one of two true states. The common prior is $\Pr(T = 2) = 0.7$, with conditional signal probabilities $\Pr(S = h \mid T = 2) = 0.8$ and $\Pr(S = h \mid T = 1) = 0.1$. The prior probability that agent i will receive a high signal is $\Pr(S_i = h) = \Pr(S_i = h \mid T = 2) \Pr(T = 2) + \Pr(S_i = h \mid T = 1) \Pr(T = 1) = 0.59$. By experiencing the product, an agent learns something about the world state. For example, following a high signal, she updates her state belief to posterior $\Pr(T = 2 \mid S_i = h) = \frac{\Pr(S_i = h \mid T = 2) \Pr(T = 2)}{\Pr(S_i = h)} = 0.95$. (The analogous update following a low signal is 0.34.) The agent also revises her belief that agent j received a high signal, with posterior signal beliefs $p(h) = \Pr(S_j = h \mid S_i = h) = 0.76$ and $p(l) = \Pr(S_j = h \mid S_i = l) = 0.34$. 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. Now assume that agent 1 and agent 2 report a low signal and a high signal, respectively. If the logarithmic rule R_l is used, the computed payments for agent 1 and agent 2 are $R_l(0.34, 1) = \ln(0.34) = -1.08$ and $R_l(0.76, 0) = \ln(0.24) = -1.43$ respectively. To ensure ex post individual rationality, the center uses a scaled version of the logarithmic rule R'_l , with an added constant corresponding to the absolute value of the lowest possible negative payment, such that $R'_l = R_l + |\ln(0.24)| = R_l + 1.43$.

There are two problems in extending the classical peer prediction method to incorporate subjective and private priors. First, it is no longer sufficient for the mechanism to only ask for signal reports because it would not be able to infer posterior signal beliefs. This could be solved by eliciting posterior signal beliefs, but then we run into a second problem which is more severe: without the prior being common knowledge, eliciting only the signal posterior *does not enable the mechanism to infer the agent's signal*. Eliciting the signal is crucial because it is used as the event that another agent shall predict.

Example 2. Consider again the setting from Example 1: if the mechanism did not know that $\Pr(T = 2) = 0.7$, $\Pr(S = h \mid T = 2) = 0.8$ and $\Pr(S = h \mid T = 1) = 0.1$, it could not infer anything from the agent reporting her signal posterior belief to be 0.34 because this could, for example, also stem from a high signal in a setting with $\Pr(T = 2) = 0.06125$ and the same conditional signal beliefs.

As a possible solution for settings with private and subjective priors, MRZ briefly discuss the possibility of a direct-revelation approach where the agents are asked to report all private information, including the prior beliefs on world state, the signal beliefs conditional on state and the signal itself. In fact, this approach is *not* strictly incentive compatible if all information is reported simultaneously.

Example 3. Consider again the setting from Example 1: if agent i observes a low signal, a truthful report of all private information including the low signal yields the same payment as a misreport of the same conditional signal beliefs but prior belief $\Pr(T = 2) = 0.06125$ together with a high signal.¹

¹Of course, if strict truthfulness is not required, the naive mechanism that asks only for a signal report and pays each agent a constant amount independent of the reported signal is a much simpler, weakly incentive compatible solution.

While MRZ do not mention this, their direct-revelation approach can be made strictly incentive compatible for admissible belief types by temporal separation. The center must ensure that the agent reports her subjective prior *before* receiving her signal. Nevertheless, this direct approach appears impractical because of its high reporting costs. Observe that in the case of 2 states and 2 signals, an agent has to report three probabilities and a signal. These reporting costs grow with the number of states, so that with 3 states and 2 signals, it would already require each agent to report five probabilities and a signal.

Thus, we believe that a different approach is required for settings with subjective priors, in order to design a mechanism that is both strictly incentive compatible, and feasible with respect to the agents' reporting costs.

4. PRIVATE-PRIOR PEER PREDICTION: THE BASIC MECHANISM

We deviate from the classical peer prediction method in two aspects. First, every agent has her own subjective belief type in regard to the world state and the way in which signals are generated given each state. Second, this belief type is private to an agent and not known by other agents or the mechanism. The difficulty comes from the second relaxation: if the mechanism knew each agent's subjective beliefs, it could still compute the possible posterior beliefs for the other agent's signal and the classical peer prediction method could be applied.

The Bayesian Truth Serum (BTS) [Prelec 2004] (for a large enough number of agents) and the Robust Bayesian Truth Serum (RBTS) [Witkowski and Parkes 2012] (for 3 or more agents) provide a solution when all agents share a *common prior*, but where this prior is not known by the mechanism. In addition to the signal report, these mechanisms ask agents to report their posterior signal beliefs and score agents on the basis of both of these reports. The following table provides a summary of the different settings:

		Prior	
		common	subjective
public		Classical Peer Prediction	(Classical Peer Prediction)
private		(Robust) Bayesian Truth Serum	this paper

In what follows we provide a first proposal, the *basic private-prior peer-prediction mechanism (BPP)*, for the setting of private and subjective priors.

4.1. The Basic Private-Prior Peer-Prediction Mechanism (BPP)

For every agent i , choose another agent $j = i + 1$ (modulo n).

- (1) Ask agent i for her prior signal belief report $y_i \in [0, 1]$ that another agent will receive a high signal.
- (2) Agent i observes signal $S_i = s_i$.
- (3) Ask agent i for her posterior signal belief report $y'_i \in [0, 1]$, with $y'_i \neq y_i$, that another agent will receive a high signal.
- (4) Infer agent i 's implicit signal report x_i by applying

$$x_i = x(y_i, y'_i) = \begin{cases} h, & \text{if } y'_i > y_i \\ l, & \text{if } y'_i < y_i \end{cases} \quad (6)$$

Agent i 's score is calculated as:

$$u_i = R(y_i, x_j) + R(y'_i, x_j), \quad (7)$$

where R is a strictly proper scoring rule, and x_j is the inferred signal report of agent j .

As we will show, the true prior signal belief and the true signal posterior cannot be the same given admissible belief types, so that $y'_i \neq y_i$ is not restrictive for a truthful agent.²

4.2. Incentive Compatibility

It is critical for the incentive properties of BPP that Step (1) happens before the agent observes signal S_i and Step (3) happens after the agent has observed signal S_i . This is the property of temporal separation. In establishing incentive compatibility, we need the following lemma:

LEMMA 4.1. *[Witkowski and Parkes 2012] For every agent i with admissible belief type, it holds that $p_i(h) > p_i > p_i(l)$.*

In words, agent i 's belief that another agent receives a high signal strictly increases from her prior belief in the event that she observes a high signal. Analogously, her belief that another agent receives a high signal strictly decreases relative to her prior belief if she observes a low signal.

In formalizing the equilibrium concept for the case of private priors, we recall that each agent has an admissible belief type $\theta_i \in \Theta$. We will not require agent i to form any beliefs about the type of another agent j . A strategy σ_i in the setting of BPP describes an agent's prior belief report $y_i \in [0, 1]$, and then for each signal $s_i \in \{0, 1\}$, an agent's posterior belief report $y'_i \in [0, 1]$. Put together, we have strategy $\sigma_i = (\sigma_i^1, \sigma_i^2)$, where $\sigma_i^1 : \Theta \rightarrow [0, 1]$ and $\sigma_i^2 : \Theta \times \{0, 1\} \rightarrow [0, 1]$. Component $\sigma_i^1(\theta_i)$ defines the agent's report y_i for every possible belief type she might have, whereas component $\sigma_i^2(\theta_i, s_i)$ defines the agent's report y'_i for every possible belief type and every possible signal.

Definition 4.2. Strategy profile $(\sigma_1, \dots, \sigma_n)$ is an *ex post subjective equilibrium* of BPP with n agents if, for every agent i , and every admissible belief type θ_i ,

$$E[R(\sigma_i^1(\theta_i), X_j)] + E[R(\sigma_i^2(\theta_i, s_i), X_j) \mid S_i = s_i] \geq E[R(\hat{y}_i, X_j)] + E[R(\hat{y}'_i, X_j) \mid S_i = s_i],$$

where $X_j = x(\sigma_j^1(\theta_j), \sigma_j^2(\theta_j, s_j))$ is the signal reported by agent j given $S_j = s_j$, and where this holds for all admissible belief types θ_j , and all belief reports $\hat{y}_i \neq \sigma_i^1(\theta_i)$ and $\hat{y}'_i \neq \sigma_i^2(\theta_i, s_i)$ (where this second report can depend on agent i 's realized signal).

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 types, and given 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. In particular, we are interested in *ex post subjective incentive compatible* mechanisms, where the truthful strategy $\sigma_i^1(\theta_i) = p_i$ and $\sigma_i^2(\theta_i, s_i) = p_i(s_i)$ is an equilibrium. That is, every agent reports her true prior signal belief $p_i = \Pr_i(S_i = h)$ and then her true posterior signal belief $p_i(s_i) = \Pr_i(S_j = h \mid S_i = s_i)$. A mechanism is *strictly* incentive compatible when the constraints defining the equilibrium hold strictly.

THEOREM 4.3. *BPP is strictly ex post subjective incentive compatible if every agent's belief type is admissible given temporal separation.*

²To keep the user interface simple, a practical deployment might allow the two reports from an agent to be equal and still pay the agent as described. In this case, the method would skip the agent's role as providing a prediction target for another agent. In the extreme case, where there is some agent i for which all other agents $j \neq i$ report $y_j = y'_j$, a practical deployment could score i against a random signal. It bears emphasis that these details do not affect the equilibrium analysis, but are all robustness issues in regard to a practical deployment.

PROOF. Assume agent $j \neq i$ is truthful. First, for agent j 's inferred signal report it holds that $x_j = x(y_j, y'_j) = s_j$ with $S_j = s_j$. To see this, verify that by Lemma 4.1 it holds that $y'_j = p_j(s_j) > p_j = y_j$ if and only if $s_j = h$, and that $y'_j = p_j(s_j) < p_j = y_j$ if and only if $s_j = l$. It follows that agent i has strict incentives to report y_i and y'_i truthfully, with respect to her subjective belief type because her score is the sum of two proper scoring rules applied to x_j and recognizing that the inference in regard to x_i does not affect agent i 's score. \square

4.3. Individual Rationality

In contrast to the classical peer prediction method, the particular choice of scoring rule matters as to whether BPP provides individual rationality. In particular, the logarithmic scoring rule R_l cannot provide individual rationality for BPP. This is because an agent's signal posterior can take values arbitrarily close to 0, and thus there is no suitable constant the mechanism could add to always make the score non-negative. Instead, we can adopt a strictly proper scoring rule that guarantees non-negative values for every possible report, so that there is no need for adding a lower bound. For example, we can adopt the *binary quadratic scoring rule* R_q , normalized to give scores between 0 and 1:

$$\begin{aligned} R_q(y, \omega = 1) &= 2y - y^2 \\ R_q(y, \omega = 0) &= 1 - y^2. \end{aligned} \tag{8}$$

PROPOSITION 4.4. [e.g., Selten 1998] *The binary quadratic scoring rule R_q is strictly proper.*

Example 4. Consider again the setting of Example 1: suppose that an agent has prior belief $\Pr_i(T = 2) = 0.7$ and conditional signal beliefs $\Pr_i(S = h \mid T = 2) = 0.8$ and $\Pr_i(S = h \mid T = 1) = 0.1$. Here, we situate the example concretely in the context of the purchase of a digital camera online. The procedure using BPP together with the quadratic scoring rule R_q is then as follows:

- (1) Agent i buys a digital camera from an e-commerce platform.
- (2) The platform asks the agent for her prior signal belief report and she truthfully reports $y_i = 0.59$.
- (3) Some days later, the agent receives and experiences the camera. She is disappointed with the picture quality, so her signal is low. She updates her signal belief to $p_i(l) = 0.34$ and truthfully reports $y'_i = 0.34$.
- (4) Another agent j buys the same camera on the platform, and also follows Steps (1) to (3), with potentially different beliefs and experiences.
- (5) Agent i is scored against agent j 's implicitly reported signal x_j . Thus, in the case where agent j was happy with the camera, agent i is paid $R_q(0.59, 1) + R_q(0.34, 1) = 2 - 0.59^2 - 0.34^2 = 1.54$.

5. PRIVATE-PRIOR PEER PREDICTION: THE SHADOW MECHANISM

In Section 4 we have seen that subjective and private priors can be handled, but with higher reporting costs compared to the classical peer prediction method. Whereas classical peer prediction requires a report of an agent's signal, BPP requires an agent to report a prior signal belief and a posterior signal belief. In this section we modify the basic mechanism so that an agent's second report is a signal report rather than a belief report. We first introduce a candidate "shadow peer-prediction" mechanism, in which the equilibrium provides strict incentives for truthful signal reports but not for truthful belief reports. Section 5.2 provides an analysis of this mechanism, identifying the

optimal deviation from the true prior belief report. By making an appeal to the *revelation principle* from mechanism design theory, we then construct the shadow private-prior peer-prediction mechanism, which is strictly incentive compatible for both reports. This is presented in Section 5.3. In Section 5.4, we present a compact variation of this mechanism, in which the payment rule consists of only a single application of the quadratic scoring rule. In Section 5.5, we show that this compact variant is individual rational and provide an illustrative example.

5.1. The Candidate Shadow Private-Prior Peer-Prediction Mechanism (Candidate SPP)

To begin, we introduce the *candidate shadow private-prior peer-prediction mechanism (Candidate SPP)* which is a building block towards a fully incentive-compatible mechanism. Candidate SPP is parametrized by $\delta > 0$:

For every agent i , choose another agent $j = i + 1$ (modulo n).

- (1) (Stage 1) Ask agent i for her prior signal belief report $y_i \in [0, 1]$ that another agent will receive a high signal.
- (2) Agent i observes signal $S_i = s_i$.
- (3) (Stage 2) Ask agent i for her *signal report* $x_i \in \{0, 1\}$, and calculate shadow posterior belief

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

where $\delta > 0$ is a parameter of the mechanism.

Agent i 's score is then

$$u_i = R_q(y_i, x_j) + R_q(y'_i, x_j), \quad (10)$$

where $R_q(\cdot, \cdot)$ is the quadratic scoring rule and x_j is the signal report of agent j .

The shadow posterior y'_i in Candidate SPP might fall outside $[0, 1]$. For this reason, we extend the definition of the quadratic scoring rule to allow belief reports outside $[0, 1]$, noting that the score for both $\omega = 0$ and $\omega = 1$ remains well defined.

5.2. Equilibrium Analysis

The subjective equilibrium concept and strict ex post subjective incentive compatibility extend in the natural way from Section 4: as in BPP, each agent is best-responding to the strategy of every other agent given knowledge of her own belief type, and with strict uncertainty on the belief types of other agents other than their admissibility. The only difference to the concepts used in Section 4 is that an agent's second report is now a signal and not a belief.

From a game-theoretic point of view, the key difference between Candidate SPP and BPP is that in Candidate SPP there is an interdependency between an agent's first and second report. Agent i 's first report y_i has an influence on the payment that she will receive for her second report x_i through its effect on y'_i . This requires careful incentive analysis. We begin with a technical lemma, slightly adapted from Friedman [1983] to accommodate out-of-bounds values.

LEMMA 5.1 (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 $|y - p|$.*

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

This property is specific to the quadratic scoring rule. In particular, the logarithmic rule does not satisfy it. This lemma does not require the set of Y to lie in $[0, 1]$, and holds for values y outside this range. This is helpful because it allows the incentives for truthful stage 2 reports to be analyzed without requiring a particular condition on $\delta > 0$ (and in particular without requiring that δ depends on the stage 1 report, which would complicate the incentive analysis).

LEMMA 5.2. *In Candidate SPP, given signal belief report y_i and realized signal s_i and assuming agent j is truthful, agent i 's optimal signal report x_i depends on her prior signal belief report y_i as follows:*

- (1) *If $y_i < p_i(s_i)$, then she has a strict preference to report $x_i = h$.*
- (2) *If $y_i = p_i(s_i)$, then she is indifferent between $x_i \in \{l, h\}$.*
- (3) *If $y_i > p_i(s_i)$, then she has a strict preference to report $x_i = l$.*

In particular, if $p_i(l) < y_i < p_i(h)$, then the truthful report $x_i := s_i$ is optimal.

PROOF. Fix belief report y_i and true signal s_i , and assume j 's signal report is truthful. From Lemma 5.1 it follows that agent i should report x_i that leads to a shadow posterior y'_i with minimal distance to agent i 's true posterior $p_i(s_i) = \Pr_i(S_j = h \mid S_i = s_i)$. Consider two cases:

- (1) *If $y_i < p_i(s_i)$ then $y_i - p_i(s_i) < 0$, and $|y_i + \delta - p_i(s_i)| < |y_i - \delta - p_i(s_i)|$ and $x_i = h$ is strictly optimal. The case of $y_i > p_i(s_i)$ is symmetric.*
- (2) *If $y_i = p_i(s_i)$ then the distance is the same for either report, and so indifference.*

This completes the proof. \square

Lemma 5.2 gives a hint as to what we're looking for— conditions on δ and an agent's subjective prior, such that $p_i(l) < y_i < p_i(h)$, and the agent's signal report is truthful.

PROPOSITION 5.3. *In Candidate SPP, if mechanism parameter $0 < \delta \leq 2(p_i(h) - p_i(l))$, agent i 's strict best response to a truthful signal report by agent j is to make belief report $y_i = p_i(1 - \delta) + \frac{\delta}{2}$ and truthful signal report $x_i = s_i$.*

PROOF. First of all, let us constrain agent i 's strategy to reporting the true signal $x_i = s_i$. Given this, the expected score for reporting y_i in stage 1 is: $U_{\text{truesignal}}(y_i) = p_i(2y_i - y_i^2) + (1 - p_i)(1 - y_i^2) + p_i(2(y_i + \delta) - (y_i + \delta)^2) + (1 - p_i)(1 - (y_i - \delta)^2)$. Taking the derivative with respect to y_i , and setting to zero, $\frac{\partial U_{\text{truesignal}}(y_i)}{\partial y_i} = 2(p_i - y_i - 2\delta p_i + \delta + p_i - y_i) = 0 \Leftrightarrow y_i = p_i + \frac{\delta(1 - 2p_i)}{2} = p_i(1 - \delta) + \frac{\delta}{2}$, and a maximum by checking second-order conditions.

It is reassuring to confirm that for $\delta \leq 2(p_i(h) - p_i(l))$, we have $p_i(l) < y_i < p_i(h)$, and thus not only a feasible report, i.e. in the $[0, 1]$ bound, but also consistency with Lemma 5.2: first, observe that for $p_i = 0.5$, it holds that $y_i = 0.5$ for any δ , so that y_i is in bounds. Second, observe that for $p_i > 0.5$, it holds that $y_i < p_i$, and for $p_i < 0.5$, it holds that $y_i > p_i$. Because of symmetry, it is then sufficient to show that $y_i < p_i(h)$ for

$p_i < 0.5$, and thus $p_i(l) < p_i < y_i < p_i(h)$. We have,

$$\begin{aligned} y_i &= p_i + \frac{\delta(1-2p_i)}{2} \leq p_i + \frac{2(p_i(h) - p_i(l))(1-2p_i)}{2} = p_i + (p_i(h) - p_i(l))(1-2p_i) \\ &= p_i + p_i(h) - 2p_i p_i(h) - p_i(l) + 2p_i p_i(l) = p_i - p_i p_i(h) - p_i(l) + p_i p_i(l) + p_i(h) \\ &\quad + p_i p_i(l) - p_i p_i(h) = p_i(h) + p_i p_i(l) - p_i p_i(h) < p_i(h). \end{aligned}$$

For the final equality, we need that $p_i = p_i p_i(h) + (1-p_i)p_i(l)$ and thus $p_i - p_i p_i(h) - p_i(l) + p_i p_i(l) = 0$, and the strict inequality follows from $p_i > 0$ and $p_i(h) > p_i(l)$, which holds for admissible belief types.

What other strategy might be better for the agent? We know from Lemma 5.2 that the only other case to consider is $y_i \leq p_i(l)$ (or symmetrically, $y_i \geq p_i(h)$), where the optimal signal report is $x_i = h$, independent of realized signal s_i . Given this, let's now constrain agent i 's strategy to always reporting h . Given this, the expected score for reporting y_i in stage 1 is $U_{\text{alwayshigh}}(y_i) = p_i(2y_i - y_i^2) + (1-p_i)(1 - y_i^2) + p_i(2(y_i + \delta) - (y_i + \delta)^2) + (1-p_i)(1 - (y_i + \delta)^2)$. Taking the derivative with respect to y_i , and setting to zero, $\frac{\partial U_{\text{alwayshigh}}(y_i)}{\partial y_i} = 2(p_i - y_i - \delta + p_i - y_i) = 0 \Leftrightarrow y_i = p_i - \frac{\delta}{2}$, and a maximum by checking the second-order conditions. However, Candidate SPP insists on $y_i \in [0, 1]$ and so for $\delta > 2p_i$ this is not feasible. Therefore, the expected utility given $y_i = p_i - \delta/2$ is an upper bound on the actual utility available when playing the "always-high" strategy in stage 2.

Continuing, we establish that the expected loss, relative to being able to report $y_i = p_i$ in stage 1 and $y'_i = p_i(s_i)$ in stage 2 is greater under the "always-high" strategy than the "true signal" strategy. The expected loss for the "true signal" strategy is:

$$\begin{aligned} L_{\text{truesignal}} &= \left(\frac{\delta(1-2p_i)}{2} \right)^2 + p_i \left(p_i(h) - \left[p_i(1-\delta) + \frac{\delta}{2} + \delta \right] \right)^2 \\ &\quad + (1-p_i) \left(p_i(l) - \left[p_i(1-\delta) + \frac{\delta}{2} - \delta \right] \right)^2 \end{aligned}$$

For the always high strategy, the expected loss is:

$$L_{\text{alwayshigh}} \geq \left(\frac{\delta}{2} \right)^2 + p_i \left(p_i(h) - \left[p_i - \frac{\delta}{2} + \delta \right] \right)^2 + (1-p_i) \left(p_i(l) - \left[p_i - \frac{\delta}{2} + \delta \right] \right)^2$$

This is a lower bound on loss, because the optimal y_i in this case may be out of the $[0, 1]$ bound and thus the agent's utility is upper-bounded by assuming $y_i = p_i - \delta/2$ is feasible. Combining, we have:

$$\begin{aligned} L_{\text{alwayshigh}} - L_{\text{truesignal}} &\geq \left(\frac{\delta}{2} \right)^2 - \left(\frac{\delta(1-2p_i)}{2} \right)^2 \\ &\quad + p_i \left(\left(p_i(h) - \left[p_i + \frac{\delta}{2} \right] \right)^2 - \left(p_i(h) - \left[p_i(1-\delta) + \frac{3\delta}{2} \right] \right)^2 \right) \\ &\quad + (1-p_i) \left(\left(p_i(l) - \left[p_i + \frac{\delta}{2} \right] \right)^2 - \left(p_i(l) - \left[p_i(1-\delta) - \frac{\delta}{2} \right] \right)^2 \right) \\ &= \delta^2 p_i(1-p_i) + \delta(1-p_i)p_i \left((-\delta + p_i(h) - p_i) + \delta(p_i - 1) + p_i(h) - p_i \right) \\ &\quad - \delta(1-p_i)(1+p_i)(p_i(l) - p_i + \delta p_i + p_i(l) - p_i) \\ &= 2\delta(1-p_i)(p_i - p_i(l) - \delta p_i + p_i(h)p_i - p_i(l)p_i). \end{aligned}$$

Since $2\delta(1 - p_i) > 0$, for $L_{\text{alwayshigh}} - L_{\text{truesignal}} > 0$ we need

$$\begin{aligned} p_i - p_i(l) - \delta p_i + p_i(h)p_i - p_i(l)p_i &\geq 0 \Leftrightarrow (p_i - p_i(l)) + p_i(p_i(h) - p_i(l)) \geq p_i\delta \\ \Leftrightarrow p_i(p_i(h) - p_i(l)) + p_i(p_i(h) - p_i(l)) &\geq p_i\delta \Leftrightarrow \delta \leq 2(p_i(h) - p_i(l)). \end{aligned}$$

This establishes that a prior signal belief report $y_i = p_i + \frac{\delta(1-2p_i)}{2}$ and the truthful signal report constitute a strict best response to a truthful signal report by agent j . \square

We immediately have the following:

THEOREM 5.4. *In a strict ex post subjective equilibrium of Candidate SPP, and given temporal separation, every agent i reports her true signal s_i if mechanism parameter $0 < \delta \leq 2(p_i(h) - p_i(l))$ for all $i \in \{1, \dots, n\}$.*

This provides partial incentive compatibility: an agent has strict incentives to report her true signal in stage 2 but should deviate to $y_i = p_i + \frac{\delta(1-2p_i)}{2}$ in stage 1. The only informational requirement on the mechanism is that it must pick a δ small enough, depending on knowledge of a valid $\epsilon > 0$ such that $\epsilon \leq p_i(h) - p_i(l)$ for all $i \in \{1, \dots, n\}$ (with $\delta \leq 2\epsilon$ being sufficient). Certainly, admissibility implies $p_i(l) < p_i < p_i(h)$ for all belief types. What is required in addition is knowledge of this minimal “degree of informativeness” of a signal, valid for every belief type in the population.

5.3. The Shadow Private-Prior Peer-Prediction Mechanism (SPP)

We can now apply the revelation principle and achieve strict incentive compatibility in regard to both the belief report and the signal report. The crucial observation is that the optimal misreport $y_i = p_i + \frac{\delta(1-2p_i)}{2}$ depends *only* on the agent’s prior signal belief p_i and parameter δ of the mechanism. It does not depend on any other aspect of the agent’s belief type. For this reason, the mechanism can simply compute y_i directly on behalf of the agents on the basis of p_i . This leads us to the *shadow private-prior peer-prediction mechanism (SPP)*:

For every agent i , choose another agent $j = i + 1$ (modulo n).

- (1) (Stage 1) Ask agent i for her prior signal belief report $y_i \in [0, 1]$ that another agent will receive a high signal.
- (2) Agent i observes signal $S_i = s_i$.
- (3) (Stage 2) Ask agent i for her signal report $x_i \in \{0, 1\}$, and calculate shadow posterior belief

$$y'_i = y'(y_i, x_i) = \begin{cases} y_i + \frac{\delta(1-2y_i)}{2} + \delta, & \text{if } x_i = 1 \\ y_i + \frac{\delta(1-2y_i)}{2} - \delta, & \text{if } x_i = 0, \end{cases} \quad (11)$$

where $\delta > 0$ is a parameter of the mechanism.

Agent i ’s score is then

$$u_i = R_q \left(y_i + \frac{\delta(1-2y_i)}{2}, x_j \right) + R_q(y'_i, x_j), \quad (12)$$

where $R_q(\cdot, \cdot)$ is the quadratic scoring rule, and x_j the report of agent j .

Observe that, for $\delta = 1$, the transformed prior belief report $y_i + \frac{\delta(1-2y_i)}{2}$ becomes $\frac{1}{2}$ which does not depend on y_i . This would mean that while agent i still cannot do better than reporting $y_i = p_i$, the prior report would not be *strictly* incentivized. Excluding this case, the following theorem then follows immediately from Proposition 5.3 and the equilibrium analysis of Candidate SPP:

THEOREM 5.5. *SPP is strictly ex post subjective incentive compatible given temporal separation, if mechanism parameter $\delta \neq 1$ and $0 < \delta \leq 2(p_i(h) - p_i(l))$ for all $i \in \{1, \dots, n\}$.*

5.4. The Compact Shadow Private-Prior Peer-Prediction Mechanism (Compact SPP)

In SPP, the signal prior report y_i affects both parts of agent i 's score while signal report x_i only affects the second part. This raises the question as to whether the second part alone could suffice for strict truth-telling incentives. That is, whether Candidate SPP's payment rule can be shortened to $R_q(y'_i, x_j)$.

In fact, this is possible. First observe that Lemma 5.2 still holds since this analysis pertained only to the score for the signal report. Given this, the outline of the analysis follows as before. Full details are deferred to the full version of the paper. First, we can derive (a) the optimal report y_i given that the agent reports her true signal s_i , and (b) the optimal report y_i given that the agent always reports a high signal (or symmetrically, always reports a low signal.) Ignoring the symmetric case of “always low”, this yields:

- (1) If $x_i = h$ (independent of s_i), the optimal belief report is $y_i = p_i - \delta$.
- (2) If $x_i = s_i$, the optimal belief report is $y_i = p_i(1 - 2\delta) + \delta$.

Considering the expected loss relative to being able to make perfect reports $p_i(l)$ or $p_i(h)$ in stage 2 (depending on the observed signal), and requiring that the loss from reporting the true signal and $y_i = p_i(1 - 2\delta) + \delta$ is strictly less than that of always reporting a high signal and $y_i = p_i - \delta$, the constraint on parameter δ is $0 < \delta \leq p_i(h) - p_i(l)$.³ Given this, and adopting the revelation principle as before, we have the following *compact shadow private-prior peer-prediction mechanism (Compact SPP)*:

For every agent i , choose another agent $j = i + 1$ (modulo n).

- (1) (Stage 1) Ask agent i for her prior signal belief report $y_i \in [0, 1]$ that another agent will receive a high signal.
- (2) Agent i observes signal $S_i = s_i$.
- (3) (Stage 2) Ask agent i for her signal report $x_i \in \{0, 1\}$, and agent i 's score is

$$u_i = R_q((1 - 2\delta)y_i + 2\delta x_i, x_j), \quad (13)$$

where $\delta > 0$ is a parameter of the mechanism, $R_q(\cdot, \cdot)$ is the quadratic scoring rule, and x_j is the signal report of agent j .

Analogous to SPP, the transformed prior report in Compact SPP is independent of y_i if $\delta = \frac{1}{2}$, so that, in that case, the prior belief report y_i would only be weakly truthful. We thus exclude $\delta = \frac{1}{2}$ in the following theorem which we state without proof:

THEOREM 5.6. *Compact SPP is strictly ex post subjective incentive compatible given temporal separation if parameter $\delta \neq \frac{1}{2}$ and $0 < \delta \leq p_i(h) - p_i(l)$ for all $i \in \{1, \dots, n\}$.*

The payment rule of Compact SPP has a very nice and intuitive interpretation: noting that it can be written as $R_q((1 - \eta)y_i + \eta x_i, x_j)$ with $\eta = 2\delta$ and observing that the expected value of x_i is p_i , then—given that agent i is truthful—the expected belief report that is applied to the quadratic scoring rule is $E[(1 - \eta)y_i + \eta x_i] = p_i$. This makes sense given that the quadratic scoring rule is strictly proper, and p_i is agent i 's best prediction for the signal report x_j at stage 1 given that agent j is truthful. The role of $\eta > 0$ is to put some “weight” on the signal report, so that the agent has an incentive to use the signal report in obtaining better accuracy in regard to her posterior signal beliefs. On the other hand, if η gets too large relative to the effect of an agent's signal on her posterior signal belief, then the agent prefers not to “adjust” her shadow posterior

³For $0 < \delta \leq p_i(h) - p_i(l)$, one can again confirm that $y_i = p_i(1 - 2\delta) + \delta$ is a feasible report. As in the analysis of SPP, it is irrelevant that $y_i = p_i - \delta$ may not be in the $[0, 1]$ range since requiring a report inside the range can only further reduce an agent's utility.

through reporting a signal in a way that depends on the observed signal. Instead, in this situation where the two possible signal posteriors are relatively close to each other, she will set $E[(1 - \eta)y_i + \eta x_i] = p_i$ by choosing $x_i = h$ (always) and $y_i = \frac{p_i - \eta}{1 - \eta} = \frac{p_i - 2\delta}{1 - 2\delta}$. By applying this to the transformation due to the direct-revelation approach, one rediscovers the optimal misreport given a high signal report $\frac{p_i - 2\delta}{1 - 2\delta}(1 - 2\delta) + \delta = p_i - \delta$.

An interesting special case is choosing $\delta = 0.5$, which is a valid parametrization for *weakly truthful* prior belief incentives when $0.5 \leq p_i(h) - p_i(l)$ and thus $p_i(h) > 0.5$ and $p_i(l) < 0.5$. In this case, the effect of belief report y_i disappears and the shadow belief report that is adopted in the scoring rule is 1 if the signal report is h and 0 if the signal report is l . It makes sense that this would retain strict incentives in regard to the signal report in this case, since $p_i(h) > 0.5$ and $p_i(l) < 0.5$ and so shadow posteriors 1 and 0 are minimizing the distances to $p_i(h)$ and $p_i(l)$, respectively.

5.5. Individual Rationality

In BPP, which utilizes two belief reports, using the normalized quadratic scoring rule R_q giving scores between 0 and 1 ensures *ex post* individual rationality. For the mechanisms from Section 5, it is no longer obvious that this still holds because there could be out-of-bound shadow posterior reports (outside of $[0, 1]$), so that the *ex post* scores may be negative. For Compact SPP and $\delta > 0.5$, this can indeed be the case. The lowest possible score is attained through the lowest possible $(1 - 2\delta)y_i + 2\delta x_i$ and $x_j = 1$, or through the highest possible $(1 - 2\delta)y_i + 2\delta x_i$ and $x_j = 0$. For $\delta > 0.5$, $(1 - 2\delta)y_i + 2\delta x_i$ is minimized by reporting $y_i = 1$ and $x_i = 0$, and maximized by reporting $y_i = 0$ and $x_i = 1$. Applied to R_q together with $x_j = 1$ and $x_j = 0$, respectively, one obtains the lowest possible score as $2(1 - 2\delta) - (1 - 2\delta)^2 = 1 - (2\delta)^2 = 1 - 4\delta^2$. Therefore, by adding $|1 - 4\delta^2| = 4\delta^2 - 1$ to every agent's score, Compact SPP can be made individual rational in the $\delta > 0.5$ case. For $\delta \leq 0.5$, however, Compact SPP's score is in between 0 and 1 without such a workaround. To see this, consider the range of possible values for $(1 - 2\delta)y_i + 2\delta x_i$: the signal report x_i is either 0 or 1, and so this becomes either $(1 - 2\delta)y_i$ or $(1 - 2\delta)y_i + 2\delta$, with the latter being strictly larger than the former. For $\delta \leq 0.5$, then $(1 - 2\delta)y_i \geq 0$ and $(1 - 2\delta)y_i + 2\delta \leq 1$ for all $y_i \in [0, 1]$. Moreover, 0 and 1 is obtained by reporting $x_i = y_i = 0$ and $x_i = y_i = 1$, respectively, so that for $\delta \leq 0.5$, Compact SPP attains scores between 0 and 1.

Example 5. Consider again the earlier example, with an agent having prior state belief $\Pr_i(T = 2) = 0.7$ and conditional signal beliefs $\Pr_i(S = h | T = 2) = 0.8$ and $\Pr_i(S = h | T = 1) = 0.1$. Consider again the context of the purchase of a digital camera online. The procedure using Compact SPP with $\delta = 0.2$, such that $\delta \leq p_i(h) - p_i(l) = 0.42$ and $\delta \leq 0.5$ (for individual rationality) is:

- (1) Agent i buys a digital camera from an e-commerce platform.
- (2) The platform asks the agent for her prior signal belief report and she truthfully reports $y_i = 0.59$.
- (3) Some days later, the agent receives and experiences the camera. She is disappointed with the picture quality, so her signal is low. She updates her signal belief to $p_i(l) = 0.34$.
- (4) In contrast to the basic mechanism, she then reports only her signal $x_i = l$.
- (5) Another agent j buys the same camera on the platform, and also follows Steps (1) to (4), with potentially different beliefs and experiences.
- (6) Agent i is scored against agent j 's reported signal x_j . Thus, in the case where agent j reported $x_j = h$, agent i is paid $R_q(0.59 + 2(0.2)(0 - 0.59), 1) = R_q(0.354, 1) = 2(0.354) - 0.354^2 = 0.583$

6. DISCUSSION: APPLICATIONS

In this section we offer some remarks in regard to considerations when making use of shadow peer prediction mechanisms in practical applications.

Information aggregation. In the classical peer prediction method, the center uses the common prior to compute the signal posteriors and publish this information. In our setting with private and subjective priors, the only objective information, i.e. the only information stemming from the world, are the signals. The center can therefore simply publish the percentage of positive signal reports, allowing each agent to incorporate this information into her own subjective prior. Note that this is common practice. Hotwire, for example, publishes the percentage of customers who have reported that they were satisfied with their stay at a given hotel. Another example is eBay, which publishes the percentage of positive reports for a given seller.

User interface (UI). The shadow peer-prediction mechanisms introduced in this paper require users to make reports about probabilistic beliefs. While it would be difficult to come up with a UI that makes reporting full distributions user friendly, we believe there are UIs that can achieve this for our setting, in which only a single probability is required. The user interface we suggest “hides” the probabilities from the users by adopting a point scale from 0 to 10. These points would directly correspond to probabilities, and allow users to interact with the system in a way they are familiar with from other online rating sites (albeit introducing a forced approximation to their reports).

Of course, the right choice of UI depends on the application, and probability reports may sometimes be feasible. For example, when booking a hotel on Expedia or Hotwire, the following question seems reasonably easy to answer: “*What is your prediction that another customer will recommend this hotel to a friend?*” Note that our shadow peer prediction mechanisms (SPP and Compact SPP) have the same type of reports as the Bayesian Truth Serum [Prelec 2004]: a “prediction report” about the experiences of other agents and an “experience report.” In a recent study with inexperienced human raters on Amazon Mechanical Turk, it has been shown that human agents can indeed report such information successfully [Shaw et al. 2011]. In fact, in their experimental comparison of different incentive schemes, the Bayesian Truth Serum elicited the responses with highest quality among all tested schemes.

7. CONCLUSION

In this paper, we have presented two incentive compatible mechanisms for the elicitation of truthful user feedback that escape the strong common knowledge assumptions of classical peer prediction. We believe that this development is of significant practical importance. The *compact private-prior peer-prediction mechanism (Compact SPP)* provides a particularly simple intuitive interpretation, is easy to analyze, and aligns incentives with truthful reporting of a belief on signals and then a signal. To the best of our knowledge, in terms of belief structure, the setting we study is the most general that has been studied in the context of peer prediction mechanisms. The theoretical analysis adopts a solution concept that is weaker than dominant strategy equilibrium but stronger than Bayes Nash equilibrium.

We believe that an important direction for future work is to design domain-specific mechanisms that allow the elimination of the belief report. For crowdsourcing applications, for example, it is oftentimes sufficient for the center to elicit information about the *relative rank* of two items. In the example of crowdsourced image tagging, the requester may have two sets of tags for an image and may want to know which to choose. By presenting agents with both sets, and asking for a report as to which of the two is better, an agent’s prior belief on signal should be uninformed (and thus 0.5) and we believe strict incentive compatibility can be achieved with just a signal report.

ACKNOWLEDGMENTS

We thank Yiling Chen, Felix Fischer and Malte Helmert for helpful discussion, and the anonymous reviewers for useful comments and feedback.

REFERENCES

- FRIEDMAN, D. 1983. Effective Scoring Rules for Probabilistic Forecasts. *Management Science* 29, 4, 447–454.
- GNEITING, T. AND RAFTERY, A. E. 2007. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association* 102, 359–378.
- JURCA, R. AND FALTINGS, B. 2006. Minimum Payments that Reward Honest Reputation Feedback. In *Proceedings of the 7th ACM Conference on Electronic Commerce (EC'06)*.
- 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. 2009. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research (JAIR)* 34, 209–253.
- 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)*.
- KALAI, E. AND LEHRER, E. 1993. Subjective equilibrium in repeated games. *Econometrica* 61, 1231–1240.
- MILLER, N., RESNICK, P., AND ZECKHAUSER, R. 2005. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science* 51, 9, 1359–1373.
- P. BATTIGALI, M. G. AND MOLINARI, M. C. 1992. Learning convergence to equilibrium in repeated strategic interactions: An introductory survey. *Ricerche Economiche* 96, 335–378.
- PRELEC, D. 2004. A Bayesian Truth Serum for Subjective Data. *Science* 306, 5695, 462–466.
- RAYKAR, V. C., YU, S., ZHAO, L. H., VALADEZ, G. H., FLORIN, C., BOGONI, L., AND MOY, L. 2010. Learning From Crowds. *Journal of Machine Learning Research (JMLR)* 11, 1297–1322.
- RUBINSTEIN, A. AND WOLINSKY, A. 1984. Rationalizable conjectural equilibrium: Between Nash and Rationalizability. *Games Econom. Behavior* 6, 299–311.
- SELTEN, R. 1998. Axiomatic Characterization of the Quadratic Scoring Rule. *Experimental Economics* 1, 43–61.
- SHAW, A. D., HORTON, J. J., AND CHEN, D. L. 2011. Designing Incentives for Inexpert Human Raters. In *Proceedings of the ACM 2011 Conference on Computer Supported Cooperative Work (CSCW '11)*. 275–284.
- WITKOWSKI, J. 2009. Eliciting Honest Reputation Feedback in a Markov Setting. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence (IJCAI'09)*.
- WITKOWSKI, J. 2010. Truthful Feedback for Sanctioning Reputation Mechanisms. In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence (UAI'10)*.
- WITKOWSKI, J. AND PARKES, D. 2012. A Robust Bayesian Truth Serum for Small Populations. In *Proceedings of the 26th AAAI Conference on Artificial Intelligence (AAAI'12)*.