



Peer Prediction with Private Beliefs

Citation

Witkowski, Jens, and David C. Parkes. 2011. Peer prediction with private beliefs. Proceedings of the Workshop on Social Computing and User Generated Content. San Jose, CA, June 5, 2011.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:9367004>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Peer Prediction with Private Beliefs

Jens Witkowski
Department of Computer Science
Albert-Ludwigs-Universität Freiburg
Freiburg, Germany
witkowski@informatik.uni-freiburg.de

David C. Parkes
School of Engineering and Applied Sciences
Harvard University
Cambridge, MA, USA
parkes@eecs.harvard.edu

ABSTRACT

Reputation mechanisms at online opinion forums, such as Amazon Reviews, elicit ratings from their users about the experiences with products of unknown quality and critically rely on these ratings being truthful. The peer prediction method by Miller, Resnick and Zeckhauser is arguably the most prominent truthful feedback mechanism in the literature. An obstacle with regard to its application are the strong common knowledge assumptions. Especially the commonly held prior belief about a product’s quality, although prevailing in economic theory, is too strict for this setting. Two issues stand out in particular: first, that different buyers hold different beliefs and, second, that the buyers’ beliefs are often unknown to the mechanism. In this paper, we develop an incentive-compatible peer prediction mechanism for these reputation settings where the buyers have private beliefs about the product’s inherent quality and the likelihood of a positive experience given a particular quality. We show how to exploit the temporal structure and truthfully elicit two reports: one before and one after the buyer’s experience with the product. The key idea is to infer the experience from the direction of the belief change and to use this direction as the event that another buyer is asked to predict.

1. INTRODUCTION

Almost every e-commerce website uses a reputation mechanism that collects ratings from their users. Those reputation mechanisms that are employed by online opinion forums, such as Amazon Reviews, are built to eliminate asymmetric information, and the objective is thus to reveal the inherent quality of the products to future customers. For a general overview of research on reputation mechanisms, we refer to the survey article of Dellarocas [1].

A common feature of reputation mechanisms is the dependency on honest buyer feedback. Most mechanisms in the literature simply *assume* that feedback is reported honestly. From a game-theoretic point this is problematic for two reasons: the first is the buyers’ motivation to participate at all. The feedback procedure requires the user to register an account, to log in and to fill out forms describing her¹ experiences. While this is time consuming and thus costly, the reported information benefits other customers but not the reporting buyer herself, so that standard economic theory predicts an under-provision of feedback. The second difficulty is honesty. External interests, i. e. biases towards dishonest reporting, come from a variety of motivations. Imag-

¹We refer to the buyers as female.

ine, for example, two companies competing for the same group of customers. Either company has an incentive to badmouth its competitor, to praise its own products or to pay the rating buyers to do so. The truthful elicitation of reputation feedback is thus crucial to incorporate into the design of a reputation mechanism.

The peer prediction method developed by Miller, Resnick and Zeckhauser [7] (henceforth, MRZ) incentivizes truthful feedback. It pays a buyer for her feedback depending on the feedback that was given about the same product by another buyer. The intuition behind this method is that the quality experiences of two buyers that have bought the same product should be “essentially” identical. That is, differences in experiences can occur but are captured in a noise parameter. Take a digital camera bought via Amazon as an example: while different customers may experience different quality due to noise in the production process, all buyers receive the identical model.

The peer prediction method has fostered much research, including the work by Jurca and Faltings [4, 5, 6] and our own work [8, 9]. To the best of our knowledge, however, it has not yet been applied. We conjecture that this is due to its strong common knowledge assumptions. In particular, it is assumed that every buyer shares the same prior belief about the product’s inherent quality and that the mechanism knows this prior.

To support these assumptions, MRZ put forth that there will be a sufficient rating history available in many situations so that the mechanism can estimate the prior probabilities. We argue that there is a logical flaw in this argument as it leaves open the question as to how the rating history itself was built: either people reported honestly without an incentive-compatible mechanism in place, then we do not need an incentive-compatible mechanism, or people reported dishonestly in which case we cannot use these reports to learn the prior. MRZ also suggest the possibility of an extension to the peer prediction method that can incorporate non-common priors and other private information. In a brief treatment, they propose a direct-revelation approach in which the buyers, in addition to their experiences, also report their prior belief on quality types and conditional signal beliefs given a quality type. While this would be incentive compatible, the proposal that we make in this paper is considerably more simple with respect to the reporting costs, and thus likely more practical.

Instead of asking for all private information, our mechanism only asks for two belief reports: one *before* the buyer receives the product and one *after*. The change of reports allows the mechanism to infer the received signal and use it to condition the other buyers’ payments. Proper incentives are provided to elicit both reports truthfully. Moreover, while we ask buyers to report probabilities, our user interface “hides” this by using a point scale from 0 to 10 that allows buyers to interact with the system in a way they are familiar with from other online rating sites.

2. BACKGROUND

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

2.1 Proper Scoring Rules

Proper scoring rules are functions that can be used to incentivize self-interested, rational agents to truthfully announce their private beliefs about the likelihood of a future event.

DEFINITION 1 (SCORING RULE). *A scoring rule R is a function $R : \mathcal{P} \times X \rightarrow \mathbb{R}$ with finite X that assigns a numerical score to each pair (p, x) , where p is a probability distribution and x is the event that eventually materializes.*

DEFINITION 2 (STRICTLY PROPER SCORING RULE). *A scoring rule is said to be proper if it leads to an agent maximizing her expected score by truthfully reporting her belief $p \in \mathcal{P}$ and strictly proper if the truthful report is the only report that maximizes the agent’s expected score.*

The procedure is as follows: first, the agent is asked for her belief report $p \in \mathcal{P}$. Second, an event $x \in X$ materializes and, third, the agent receives payment $R(p, x)$. An example for such a strictly proper scoring rule is the logarithmic rule that pays the agent $\log(p_x)$ where p_x is the belief that the agent assigned to the materialized event x . While the logarithmic rule always yields a negative score, there exist proper scoring rules that guarantee non-negative scores. Moreover, 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, we refer to the article by Gneiting and Raftery [2].

2.2 The Peer Prediction Method

In the peer prediction method [7], a group of rational, self-interested buyers (henceforth *agents*) experiences the same product or service. The quality of the product or service (henceforth its *type*) is unknown and never revealed. Type θ is drawn out of a finite set of possible types $\Theta = \{\theta_1, \dots, \theta_T\}$ with $T \geq 2$ and, once determined by “nature”, it is fixed. All agents share a common prior belief $p(\theta_t) = Pr(\theta = \theta_t)$ that the product is of type θ_t with $\sum_t^T p(\theta_t) = 1$ and $p(\theta_t) > 0$ for all $\theta_t \in \Theta$.

The quality observations by the agents are noisy, so that, after experiencing the product, an agent does not know with certainty the product’s actual type. Instead, she privately receives a signal drawn out of a set of signals: $S = \{s_1, \dots, s_M\}$. Let s^i denote the signal received by agent i and let $f(s_m|\theta_t) = Pr(s^i = s_m|\theta = \theta_t)$ be the probability that agent i receives the signal $s_m \in S$ if the product

is of type $\theta_t \in \Theta$. The signal observations again constitute a probability distribution, i.e. for all $\theta_t \in \Theta$ we have: $\sum_{m=1}^M f(s_m|\theta_t) = 1$. It is assumed that different types generate different conditional signal distributions and that all $f(s_m|\theta)$ are positive and common knowledge. In Section 3, in addition to dropping the assumption of a common prior, we also drop this assumption of a common-knowledge conditional signal probability $f(\cdot|\cdot)$.

The idea behind the peer prediction method is best explained with a simple example: say, there are two types, $G = \theta_2$ and $B = \theta_1$, and two signals, $s_2 = h$ for high and $s_1 = l$ for low. The type prior is $p(G) = 0.7$, and the conditional signal probabilities are $f(h|G) = 0.8$ and $f(h|B) = 0.1$. Note that $p(G)$, $f(h|G)$ and $f(h|B)$ are common knowledge. When an agent experiences the product, she receives either a high or a low signal and the mechanism’s goal is to elicit this signal truthfully. While the product’s true type is never revealed, the mechanism can compare the signal report of one agent with another agent’s signal report about the same product to create the right incentives. We refer to these agents as agent i and agent $3 - i$ (note that if $i = 1$, it holds that $3 - i = 2$ and vice versa). For the mechanism to be truthful, agent i ’s best response to a truthful report by agent $3 - i$ must also be a truthful report. The report of agent $3 - i$ will play the role of the event in a proper scoring rule.

The prior probability that agent i will receive a high signal, $Pr(s^i = h) = f(h|G) \cdot p(G) + f(h|B) \cdot p(B) = 0.59$, can be computed because all elements of the equation are common knowledge. Furthermore observe that, by experiencing the product, agent i learns something about the product’s type. For example, following a high signal, she updates her prior type belief $p(G) = 0.7$ to the type posterior $Pr(G|s^i = h) = \frac{f(h|G) \cdot p(G)}{Pr(s^i = h)} = 0.95$. (The analogous update following a low signal is 0.34.) Via this belief update, agent i also learns something about the signal of agent $3 - i$, and she can compute posterior signal beliefs $Pr(s^{3-i} = h|s^i = h) = 0.76$ and $Pr(s^{3-i} = h|s^i = l) = 0.34$. We know that proper scoring rules can be used to elicit probabilistic beliefs about an uncertain event. In our setting, this event that agent i shall predict is the signal report of agent $3 - i$. The peer prediction method takes advantage of the setting’s common knowledge and, instead of asking the agent for her entire belief, offers her only two reports, namely a high and a low report. Then, if the agent reports high, the center knows that this corresponds to a belief of 0.76 and computes the proper scoring rule with this number, with the analogous holding true for a low report.

3. PRIVATE-BELIEF PEER PREDICTION

We relax the assumption of the peer prediction method in two aspects. First, every agent has her own subjective belief with respect to the size of the (hidden) type set, the type priors and the conditional signal beliefs. Second, these beliefs are private, i.e. only known to the respective agent. Observe that the difficulty comes primarily from the second relaxation: if the center knew each agent’s subjective beliefs, it could still compute the possible posterior beliefs for the other agent’s signal. This is impossible if the subjective beliefs are also private.

3.1 Straw Man Solutions

We give an intuition as to why the peer prediction method cannot be readily extended to incorporate private beliefs in a practical way. There are two problems: first, it is no longer sufficient for the mechanism to only ask for signal reports because it does not know the individual agent’s signal posteriors that are required for the proper scoring rule. This could be solved by allowing for entire belief reports, but then we run into a second problem which is more severe: without the type priors and signal conditionals being common knowledge, eliciting only the signal posterior does not enable the center to infer the observed signal. Eliciting the signal is crucial because it is used as the event that another agent shall predict. Consider again our example numbers from Section 2.2: if the center did not know that $p(G) = 0.7$, $f(h|G) = 0.8$ and $f(h|B) = 0.1$, it could not infer anything from the agent reporting her signal posterior belief to be 0.34 since this could, for example, also stem from a high signal in a setting with $p(G) = 0.06125$ and the same signal conditionals. Another straw man solution is to ask the agents for both a prior belief on the type or the signal and the signal itself. However, it is then unclear how to incentivize the agents to report their true signal since the center can no longer compute the corresponding posterior without knowing the conditional signal beliefs.

MRZ suggest a direct-revelation approach where the agents are asked to report all private information, including the prior beliefs on types, the signal beliefs conditional on types and the signal itself. While this approach does not work if all information is reported simultaneously, it is truthful if the center ensures that the agent reports her private beliefs before experiencing the product. However, we believe that this approach is infeasible due to its high reporting costs. Note that even in the smallest possible example with only two types and two signals, the agent has to report three distributions and a signal. Moreover, these reporting costs grow in the number of types and signals, so that for the second smallest setting with three types and two signals, it would already require each agent to report five distributions and a signal.

This means that a different approach is required for settings with private information in order to design a mechanism that is both truthful and feasible with respect to the agents’ reporting costs.

3.2 The Setting

In our private-belief setting, the signal set $S = \{s_1, s_2\} = \{l, h\}$ is binary, and assumed to be common knowledge. This assumption is crucial for the implicit way that we elicit the signal reports. The true size of the type set, however, is unknown and every agent holds a private belief about it. We denote agent i ’s belief about the size of the type set by T^i with $T^i \geq 2$. In addition, both agent i ’s type prior and her conditional signal probabilities are also private information and denoted by $p^i(\theta_t)$ and $f^i(h|\theta_t)$, respectively. It is assumed common knowledge that the higher the type, the higher the probability for a high signal, i.e. $f^i(h|\theta_{t+1}) > f^i(h|\theta_t)$ for all $t \in \{1, \dots, T^i - 1\}$ and for all i . We call $(T^i, p^i(\cdot), f^i(h|\cdot))$ the *model* of agent i . Different individual beliefs about a product’s quality can for example arise if one agent has read reviews about the product on

other rating websites which—depending on the trustworthiness of these sites—may lead to better or worse estimates. It is important to note that, independent of the individual models, there is a true model $(T, p(\cdot), f(h|\cdot))$ that is underlying the signals received by agents, which is unknown to both the center and the agents. An agent’s model is thus subjective, and possibly wrong, while her signal observation is correct by definition since it stems from the true model. It is assumed that agents are risk-neutral and maximize their expected payment.

3.3 The Incentive Scheme

When an agent observes a signal, she updates her type and signal beliefs according to her model. The equations are essentially the same as those that are used in the peer prediction method. For reasons of clarity, we abbreviate the belief of agent i about the posterior probability that agent $3 - i$ received s_k , given i received s_j , as in Equation 1:

$$g^i(s_k|s_j) = Pr^i(s^{3-i} = s_k | s^i = s_j). \quad (1)$$

We can compute $g^i(s_k|s_j)$ for every signal combination by transforming Equation 1 until we are left with parameters of the agent’s model, i.e. about which she holds beliefs. The first step is to expand the conditional probability into a summation:

$$g^i(s_k|s_j) = \sum_{t=1}^{T^i} f^i(s_k|\theta_t) \cdot Pr^i(\theta = \theta_t | s^i = s_j). \quad (2)$$

Applying Bayes’ rule to the second part of the summation in Equation 2 yields:

$$Pr^i(\theta = \theta_t | s^i = s_j) = \frac{f^i(s_j|\theta_t) \cdot p^i(\theta_t)}{Pr^i(s^i = s_j)}. \quad (3)$$

The denominator of Equation 3 is the prior signal belief, and can be computed with Equation 4:

$$Pr^i(s^i = s_j) = \sum_{t=1}^{T^i} f^i(s_j|\theta_t) \cdot p^i(\theta_t). \quad (4)$$

With these, we have all necessary calculations to compute agent i ’s posterior belief about agent $3 - i$ ’s signal $g^i(s_k|s_j)$ for all $s_k, s_j \in S$.

In our setting, the reputation mechanism is situated at a trusted intermediary (e.g., at Amazon) that ensures that an agent can only give feedback about a product that she has bought. The mechanism thus also knows *when* an agent bought a product and we propose to elicit two signal belief reports: one before and one after the signal observation. If we can ensure that both reports are honest, we can infer the signal through the belief change or, intuitively, “what has been learned” by experiencing the product. To simplify notation, we overload g^i and denote the individual prior signal belief by $g^i(h) = Pr^i(s^{3-i} = h) = Pr^i(s^i = h)$ and the individual posterior signal belief by $g^i(h|s_j) = Pr^i(s^{3-i} = h | s^i = s_j)$. The following observation follows from the fact that $f^i(h|\theta_{t+1}) > f^i(h|\theta_t)$ for all $t \in \{1, \dots, T - 1\}$.

PROPOSITION 1. *For every agent i , it holds that $g^i(h|h) > g^i(h) > g^i(h|l)$.*

Agent i 's two belief reports are denoted by r_1^i and r_2^i , for the prior signal belief and the posterior signal belief, respectively. Agent i 's implicit signal report, r^i , can then be inferred in the following way:

$$r^i = \begin{cases} h, & \text{if } r_2^i > r_1^i \\ l, & \text{if } r_2^i < r_1^i \end{cases} \quad (5)$$

Note that Proposition 1 states that $g^i(h) \neq g^i(h|s_j)$. Nevertheless, to keep the user interface as simple as possible, we still allow the two reports to be equal and also pay the agent as described subsequently. We do not, however, use her reports to score another agent but randomly choose another agent.

Now that we have explained how to infer an implicit signal report from the difference in signal belief reports, we can proceed by assuming honest belief reports by agent $3 - i$ and designing strict incentives for honest belief reports by agent i . Since the signal observations themselves are not subjective but stem from the true model, we can simply use the inferred signal of agent $3 - i$ to condition the payment to agent i using a proper scoring rule: assume that agent $3 - i$ honestly reports her beliefs, i.e. $r_1^{3-i} = g^{3-i}(h)$ and $r_2^{3-i} = g^{3-i}(h|s_k)$, so that we can infer the honest signal $s_k = s^{3-i} = r^{3-i}$ by applying Equation 5. Agent i is then simply paid according to a proper scoring rule for *each* of her two reports:

$$R(r_1^i, r^{3-i}) + R(r_2^i, r^{3-i}).$$

Putting this all together, the mechanism's procedure is as follows:

1. Agent i holds prior belief $g^i(h)$ about her own and agent $3 - i$'s signal. She privately reports belief r_1^i .
2. Agent i experiences the product, i.e. receives a signal $s^i = s_j$. She updates her signal belief to $g^i(h|s_j)$ and privately reports r_2^i .
3. The analogous sequence of steps for agent $3 - i$: she reports a signal prior r_1^{3-i} , receives a signal $s^{3-i} = s_k$ and reports a signal posterior r_2^{3-i} .
4. Agent i 's reports are scored against $3 - i$'s inferred signal r^{3-i} and agent i receives $R(r_1^i, r^{3-i}) + R(r_2^i, r^{3-i})$.

THEOREM 2. *The mechanism described in Section 3.3 is perfect-Bayesian incentive compatible.*

We give an informal sketch of the proof. Consider agent i 's reporting decision before experiencing the product. Her knowledge is reflected by her model with which she can compute her subjective belief of a high signal (Equation 4). Note that the first report of an agent does not influence the payment for the second report of an agent. This simplifies the analysis, in that the (strict) best response condition then implies sequential rationality, so that the former is sufficient for incentive compatibility. Since she will be paid according to a strictly proper scoring rule, her expected payment is uniquely maximized through an honest report of $r_1^i = g^i(h)$. Remember that the subjective signal beliefs are valid for

both her own signal and the signals of any other agent drawing from the same true model. The argument for the second report is analogous: once agent i receives a signal $s^i = s_j$, she updates her subjective signal belief and—given that she is paid according to a strictly proper scoring rule—uniquely maximizes her expected payment through an honest report of $r_2^i = g^i(h|s_j)$.

In contrast to the original peer prediction method, individual rationality cannot be achieved with every proper scoring rule. The logarithmic rule, for example, can be used for the original peer prediction method but not for our private-belief mechanism. The logarithm for any number in the interval between 0 and 1 is never positive. In the original peer prediction method, the center knows all M possible signal posteriors and their respective logarithm. There, it is thus possible to add a constant to every payment such that the logarithmic rule plus this constant is non-negative for any possible signal report. In our setting with private models, however, an agent's signal posterior can take values arbitrarily close to 0. This means that there is no lower bound for the possible scores, and this makes it impossible to find a suitable constant to add.

Fortunately, there are strictly proper scoring rules that guarantee non-negative values for any possible report. An example for such a rule is the binary quadratic rule with an added constant of 1 and scaled by 0.5:

$$\begin{aligned} R(p, r^{3-i} = h) &= 2p - p^2 \\ R(p, r^{3-i} = l) &= 1 - p^2 \end{aligned} \quad (6)$$

For example, if agent i reports belief $p = 0.8$ for a high signal, she will receive 0.96 if agent $3 - i$ reports high and 0.36 if agent $3 - i$ reports low. Since every agent is guaranteed a non-negative payment, the private-belief peer prediction mechanism with the transformed proper scoring rule from Equation 6 is ex-post individually rational.

Note: A property of our mechanism is that buyers have to report probabilities. The user interface, however, "hides" these by using a point scale from 0 to 10. These points directly correspond to probabilities but instead of asking for probability reports, it allows buyers to interact with the system in a way they are familiar with from other online rating sites.

3.4 Example

In this section, we briefly exemplify the mechanism with an agent buying a digital camera from Amazon. For reasons of simplicity, we use the same numbers as in the example of Section 2.2, so imagine the agent believes that there are two types, $G = \theta_2$ and $B = \theta_1$. Her prior type belief is $p^i(G) = 0.7$, and her conditional signal beliefs are $f^i(h|G) = 0.8$ and $f^i(h|B) = 0.1$. The procedure using the transformed quadratic scoring rule from Equation 6 is then as follows:

1. The agent buys a digital camera from Amazon.
2. Amazon asks for her first signal belief report r_1^i and she truthfully reports 5.9 points, which the mechanism interprets as $g^i(h) = 0.59$.

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 $g^i(h|l) = 0.34$ and truthfully reports 3.4 points.
4. Another agent buys the same camera from Amazon and also follows Steps 1 to 3, with potentially different beliefs and experiences.
5. Our agent is scored against the other agent’s implicitly reported signal. Imagine that the other agent was happy with the camera, then our agent is paid $R(0.59, h) + R(0.34, h) = 2 - 0.59^2 - 0.34^2 = 1.54$.

3.5 Information Aggregation

Left to discuss is how the elicited signals can be aggregated, i.e. how the elicited information is fused into a joint rating. In the standard peer prediction method, one can use the true model to compute the true signal posteriors. That is, using Bayesian updating, the center can publish its belief that an agent will receive a high signal, i.e. have a positive experience with the product.

In our setting there are two natural choices. Even without beliefs, since the signals are objective, publishing the percentage of high signals of all elicited signals gives us a first approximation. Alternatively, the center could maintain its own beliefs, that it uses only for the aggregation.

4. CONCLUSION

In this paper, we have presented a mechanism for eliciting truthful reputation feedback that is in the spirit of the peer prediction method but does not rely on the same strong common knowledge assumptions. We believe that this development is of significant practical importance. In addition to truthful feedback elicitation on online opinion forums such as Amazon Reviews [e.g., 3, 4, 5, 6], the mechanism described in this paper can also be applied to online auction sites, such as eBay [e.g., 9]. In particular, we see a good fit with our work on escrow mechanisms [10], which establish trust and cooperation between traders on online auction sites without a history of feedback reports, and for an intermediary without distributional information on buyer valuations. Introducing the peer prediction method from this paper further enhances the practicality of these escrow mechanisms by removing the need for common knowledge amongst buyers about seller types. For future work, we see several exciting directions:

Learning the True Model

So far, we propose to simply publish the percentage of high signals that were implicitly reported by the agents. As mentioned earlier, an alternative is that the center itself has a model that it uses to aggregate information. Instead of assuming that the center has its own model, the center can learn this from the rich information that each agent provides. At this point, we use the two signal belief reports that an agent provides to infer a binary signal. Interestingly, the score that an agent receives for her two reports depends on her exact reports and not only the inferred signal, with higher payments corresponding to better information. In this respect, our mechanism can itself be regarded

as a scoring rule that incentivizes agents to be truthful but also estimates their private models. We plan to use this latter property to learn an agent’s private model.

System-Wide Perspective

In the setting of Section 3, we consider each product in isolation, in that the report of an agent for one product does not change the publication of reports for any other product. In reality, some agents will inherently have better information than others. This can be due to some of them studying many product reviews before buying a product or, alternatively, because they are domain experts, i.e. they have a lot of experience with similar products. In future work, we will take a system-wide perspective and take into consideration these inherent differences in agent characteristics. Combining this with a more sophisticated information aggregation scheme, this will improve the forecasting abilities of our mechanism which is especially useful for event products, such as concerts or festivals. Due to the one-shot nature of these settings, there cannot be any ratings from former customers. It is thus crucial to use the first of the two reports to already estimate the event’s quality beforehand. This is possible if there is a correlation between the accuracies of an agent’s reports for different products.

5. ACKNOWLEDGEMENTS

We thank Yiling Chen and Malte Helmert for helpful discussion. The first author is also grateful for financial support through PhD fellowships from the Landesgraduiertenförderung Baden-Württemberg and the German Academic Exchange Service.

6. REFERENCES

- [1] Chrysanthos Dellarocas. Reputation Mechanisms. In Terry Hendershott, editor, *Handbook on Information Systems and Economics*. Elsevier Publishing, 2006.
- [2] Tilmann Gneiting and Adrian E. Raftery. Strictly Proper Scoring Rules, Prediction, and Estimation. *Journal of the American Statistical Association*, 102:359–378, 2007.
- [3] Radu Jurca and Boi Faltings. Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms. In *Proc. of the 1st International Workshop on Internet and Network Economics (WINE’05)*, 2005.
- [4] Radu Jurca and Boi Faltings. Minimum Payments that Reward Honest Reputation Feedback. In *Proc. of the 7th ACM Conference on Electronic Commerce (EC’06)*, 2006.
- [5] Radu Jurca and Boi Faltings. Robust Incentive-Compatible Feedback Payments. In *Trust, Reputation and Security: Theories and Practice*, volume 4452 of *LNAI*, pages 204–218. Springer-Verlag, 2007.
- [6] Radu Jurca and Boi Faltings. Mechanisms for Making Crowds Truthful. *Journal of Artificial Intelligence Research (JAIR)*, 34:209–253, 2009.
- [7] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. *Management Science*, 51(9):1359–1373, 2005.
- [8] Jens Witkowski. Eliciting Honest Reputation Feedback in a Markov Setting. In *Proc. of the 21st International Joint Conference on Artificial Intelligence (IJCAI’09)*, 2009.
- [9] Jens Witkowski. Truthful Feedback for Sanctioning Reputation Mechanisms. In *Proc. of the 26th Conference on Uncertainty in Artificial Intelligence (UAI’10)*, 2010.
- [10] Jens Witkowski, Sven Seuken, and David Parkes. Incentive-Compatible Escrow Mechanisms. In *Proc. of the 25th AAAI Conference on Artificial Intelligence (AAAI’11)*, 2011.