Using CHI-Scores to Reward Honest Feedback from Repeated Interactions *

Radu Jurca
Ecole Polytechnique Fédérale de Lausanne
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
radu.jurca@epfl.ch

Boi Faltings
Ecole Polytechnique Fédérale de Lausanne
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
boi.faltings@epfl.ch

ABSTRACT

Online communities increasingly rely on reputation information to foster cooperation and deter cheating. As rational agents can often benefit from misreporting their observations, explicit incentives must be created to reward honest feedback. Reputation side-payments (e.g., agents get paid for submitting feedback) can be designed to make truth-telling optimal. In this paper, we present a new side-payment scheme adapted for settings where agents repeatedly submit feedback. We rate the *feedback set* of an agent, rather than individual reports. The *CHI-Score* of the feedback set is computed based on a Chi-square independence test that assesses the correlation between the agent's feedback and the feedback submitted by the rest of the community. The mechanism has intuitive appeal and generates significantly lower costs than existing incentive-compatible reporting mechanisms.

Categories and Subject Descriptors

I.2.11 [**Artificial Intelligence**]: Multiagent systems

General Terms

Algorithms, Design, Economics

Keywords

reputation mechanisms, honest reporting

1. INTRODUCTION

Reputation mechanisms emerge as a promising solution for building trust in online communities lacking objective monitoring and enforcing authorities. The experience of previous interactions (i.e., reputation information) is an indicator of future performance; when feedback is shared within the community, agents can accurately identify malicious partners and thus protect themselves against cheating. The *Feedback Forum* of eBay is a notorious example of how reputation mechanisms make business possible in a decentralized environment where two agents have few a priori reasons to trust each other.

In a service-oriented market, reputation mechanisms can do more than just identify and isolate malicious providers. Perfect service is impossible, or prohibitively expensive to provide. Occasional failures happen against the best intention of service providers; excluding such providers from the market would be unfair and inefficient. Instead, reputation information can be used to approximate the real quality of service (QoS) delivered by a provider, and consequently, to dynamically scale the price paid for service [10]. Such a flexible pricing scheme allows service providers with different characteristics to reach the best compromise between delivered quality (hence revenue) and costs. Jurca and Faltings [11] show how reputation information can become part of the Service Level Agreements (SLAs) and efficiently substitute external monitoring of Quality of Service (QoS). They prove that such *reputation-based* SLAs make it rational for service providers to keep their promises, leading to an efficient and trustworthy market.

In this paper, we address a similar scenario. Autonomous agents (clients) contract services from the different providers available in the market. The contract (i.e., the SLA) specifies the price of service, the promised QoS and a monetary penalty, paid back to the agent, when the delivered QoS is lower than promised (see [2, 11] for a framework and detailed description of such interactions). Assuming that all clients in a given group are treated equally within a certain period of time, their aggregated feedback estimates the real QoS delivered to them as a group. Reputation information (i.e., aggregated feedback) can therefore replace external monitoring of QoS, and serve to compute the penalty due by the providers.

This leaves the problem of obtaining true feedback from self-interested agents. Reporting the truth is not optimal: by submitting false negative reports, the agents can decrease the reputation of the provider and thus decrease the total price (i.e., the price specified by the SLA minus the reputation-based penalty) paid for service. Misreporting incentives can be counterbalanced by explicit rewards for honest feedback. The reputation mechanism can use side-payments to pay for the agents' reports. As no objective information is available in the market, the *reputation* side-payments are computed by correlating one agent's report

*First author is a student. We thank the anonymous reviewers for helpful remarks and comments.

AAMAS'06 May 8–12 2006, Hakodate, Hokkaido, Japan.

with reports submitted by peers. When carefully designed, these payments make it optimal for every agent to reveal the true feedback. Moreover, they can be scaled such that reporting costs or other external benefits obtained from lying (e.g., a price decrease) are accounted for. [12, 11].

However, existing side-payment schemes assume that every agent interacts only once with a service provider (in a given time period), and can therefore submit only one feedback. When a client interacts several times with the same provider, coordinated lying in multiple feedback reports can still be profitable.

Take, for example, a web service providing stock quotes. The client requests the price of a specified stock at a given time: e.g., *"Request the price of an IdealCompany share at 11:00AM on Oct 18, 2005"*. The provider promises an answer within a certain deadline: e.g., *information is delivered within 1 minute from the request time* (i.e., before 11:01 AM for the previous request). The daily SLA specifies (a) the cost of a service request, (b) the promised QoS (i.e., the probability that the request is answered before the deadline) and (c) a penalty if the delivered QoS is lower than promised, for the considered day.

One agent a that constantly monitors (e.g., once every ten minutes) the price of a share, will issue $6 \cdot 8 = 48$ service requests during one 8-working-hours day. This agent is therefore entitled to submit 48 binary feedback reports by the end of the day, where every report signals that one of the requested answers was received or not before the deadline. By misreporting only one feedback (negative instead of positive; the other 47 reports remain unchanged) the agent expects:

- to lose amount ΔSP due to the side-payment for that report;
- to decrease the price of every service invocation by Δp

Existing side-payment schemes make sure that ΔSP is always greater than Δp; however, the total price decrease of agent a is $48\Delta p$, as the price decrease also applies to the other 47 service invocations where a reported true feedback. By a similar argument, a will find it profitable to misreport all other positive reports, and she will end up reporting a clean set of 48 negative reports. Eliminating the lying incentives of agent a requires $\Delta SP \geq 48\Delta p$, which translates into a 480% increase of the original reputation side-payments.[1] The cost of reputation management simply becomes too big.

In this paper we present a new mechanism for rewarding truth-telling in settings where agents have several interactions and can submit several reports within a given time-window. The first idea behind our mechanism is to group all reports of an agent from a given time period into a *feedback set*, and pay the agent only once based on her feedback set. Second, we introduce *CHI-Scores* as a basis for computing the side-payments. The CHI-Score for a feedback set is computed based on a *Chi-square independence test* [18] applied to the feedback set itself, and the set containing the feedback of all other agents. The CHI-Score reflects the correlation between the reports of one agent and those of the entire group. Third, we use CHI-Scores not only to pay the reporter, but also to weigh different feedback sets when computing the reputation of the provider. Finally, we explicitly model external lying incentives (e.g., decrease in price) to make sure that the side-payments based on CHI-Scores makes honest reporting optimal.

Like other agreement-based (between the considered report and the general opinion) side-payment schemes (see Section 2) the mechanism we propose is not incentive compatible in general. Truthful reporting is the optimal strategy only when an agent's beliefs regarding the reports of other agents are entirely determined by her direct experience (i.e., the agent does not have any prior information). This is a common assumption for software agents in multi-agent systems and therefore we particulary address our work to such environments. Moreover, the scheme is intuitive, computationally cheap, robust against small numbers of irrational reporters, and generates significantly lower costs than existing schemes (Section 4).

Section 2 discusses related work, Section 3 formally introduces the setting and our mechanism. An experimental evaluation of the mechanism is presented in Section 4, followed by future work and the conclusion.

2. RELATED WORK

The contribution of reputation information in dynamically scaling the prices of goods and services has been established both theoretically and experimentally. Dellarocas [3] studies an "eBay-like" market where sellers can strategically choose whether to ship or not the good purchased by a buyer. He proves the existence of an equilibrium where the reputation of the seller from the last N transactions determines the present price of the good. Experimental studies of eBay transactions [5, 15] confirm that the reputation of a seller affects the selling price of the auctioned goods.

Jurca and Faltings extend the analysis of reputation-based pricing to service oriented environments where providers with different capabilities coexist [10, 11]. Their results are based on two observations. First, all clients in a given group are probably treated equally[2]. Second, service providers cannot change the quality of the delivered service infinitely often. Therefore, the reputation information obtained by aggregating the feedback from the immediate past offers an accurate estimation of the real QoS delivered by the provider. In [10], client's pay the price corresponding to yesterday's reputation. In [11], a reputation-based SLA allows clients to be refunded at the end of a given time period if the provider's reputation is lower than promised. Both mechanism promote trustworthy behavior, as service providers have the incentive to deliver the promised QoS.

The question of honest reporting has been extensively addressed by Miller et al. in [12]. The mechanism uses the feedback of a future client (called the *rater*) to rate (and compute the payment for) a submitted report. The present report is first used to update a probability distribution for the report of the rater. The payment is then computed by a scoring rule[3] that assesses how the actual report of the rater relates to the probability distribution induced by the rated report. Truthful reporting is a Nash equilibrium in all contexts: it is always optimal for an agent to report the truth

[1] for one given side-payment scheme, multiplication by a constant is the only method for increasing the side-payment loss due to misreporting

[2] note that providers can have several customer groups that are treated differently: e.g., silver/gold/platinium clients

[3] see [1] for a detailed description of scoring rules. The logarithmic function is an example of a scoring rule

given that all other agents report the truth. Moreover, the side-payments can be scaled to account for reporting costs and external benefits obtained from cheating.

Prelec [14] considers a similar approach. His mechanism asks every agent two questions: the first, regarding the observed feedback, the second regarding the feedback other agents are likely to submit. Agents are rewarded if the feedback they report is "unexpectedly common" compared with the overall prediction made by the group about the feedback likely to occur. This *information scoring* makes truthful reporting pareto optimal.

[12] discusses the danger associated with using *agreement* for computing side-payments. When the report of an agent is paid according to how well it "agrees" with the general opinion, the optimal reporting strategy is to adhere to the expected majority. Take an eBay buyer cheated by a seller with a large positive *feedback score*. Despite the negative experience, the buyer continues to believe that most of the future buyers will experience good transactions. Therefore, the report that maximizes the agreement with future feedback is the false positive one.

Agreement-based payments can therefore reward honest feedback only when the agent's expectation of future transactions coincides with her direct experience. In the case of one interaction per agent, this is possible when the outcome of successive interactions are correlated. [7] documents such a scenario where the behavior of service providers is modeled by a Markov chain satisfying certain conditions.

When agents get feedback from several direct interactions, agreement-based payment does encourage truth-telling as long as agents do not have external, trusted third-party information. This assumption is common for multi-agent systems, and has been extensively used in social networks to filter out malicious reports. In TRAVOS [16], the weights of peer reports are computed based on a Bayesian model assessing the agreement between previous recommendations and the direct experience. In CREDENCE [17] malicious reporters are filtered based on a distance measure between the agent's own observations and those of the reporter. Similar techniques are described in [20, 6, 19].

A different approach to honest feedback elicitation is presented in [8] and [13]. They rate feedback for each transaction separately, independent of other transactions. The mechanism punishes both the client and the provider if they have conflicting reports regarding the outcome of the transaction in question. [8] proves that (in any equilibrium) the number of false reports collected by such a mechanism is limited by an upper bound. [13] show an implementation of the mechanism for the Peer-to-Peer environment.

3. USING CHI-SCORES TO REWARD HONEST FEEDBACK

We propose a new side-payment scheme to explicitly reward honest feedback, based on the following ideas: first, feedback reports submitted by one agent about the same provider (within a relevant time window) are treated unitarily as one feedback set of the agent. Second, each feedback set is rated against the set of all other relevant reports available in the market for the considered time period. The CHI-Score assigned to a feedback set is computed using a Chi-square independence test applied to the two sets: the rated feedback set and the reference set containing all other

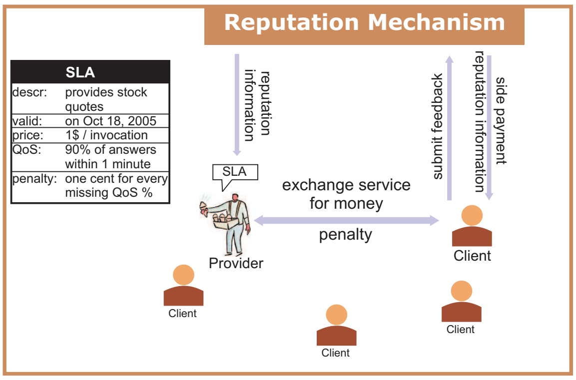

Figure 1: An online market of services.

available feedback reports. Finally, the CHI-Score is used (a) to compute the monetary reward paid to the reporting agent, and (b) to weigh the feedback when computing the reputation information. The external benefits from lying (e.g., price cuts) are factored in so that truth-telling is a Nash equilibrium.

3.1 The Scenario

We consider an online market pictured in Figure 1 where service providers repeatedly offer the same service (having the same functional interface) to the interested clients (autonomous agents), in exchange for money. Agents choose the desired providers based on the Service Level Agreement (SLA) advertised by the latter. The SLA specifies:

- the description of the service
- a period of validity,
- the quality parameters of the delivered service (i.e., the QoS),
- the price for every service invocation,
- the monetary penalty to be refunded when the actual QoS is lower than the advertised one.

A precise definition of such SLAs is presented in [11]. A practical framework supporting such interactions is documented in [2].

We assume that clients have two degrees of satisfaction: they perceive either *high* or *low* quality service. For example, high quality service means that the client obtains the answer before the deadline. QoS can therefore be defined as the probability that one service request is answered successfully (i.e., the client perceives high quality).

The QoS actually delivered to clients depends on the fixed infrastructure (e.g., computing power, bandwidth, etc) and on the strategic choices of the service provider (e.g., number of accepted requests, allocation of resources, etc). The provider can anticipate the implications of his[4] decisions on the perceptions of the clients; therefore we say that the provider can choose (within the limits imposed by the infrastructure) the QoS he is actually going to deliver.

[4]as a convention, we will refer to the provider as "he" and to the client as "she"

A service provider cannot change his decisions infinitely often. We split time into equidistant periods (indexed according to t) such that all services delivered within one period will share the same QoS. We also assume that a large enough group of clients will request the service in any period, and that their observations will be independently identically distributed according to the delivered QoS. Note that a provider can have several customers groups (i.e., silver/gold/platinium) as long as all clients in a given group are treated identically within a given period of time.

An independent reputation mechanism requests binary feedback (1 for a high quality service, 0 otherwise) about every interaction. The reputation of a provider is computed as the average satisfaction rate of the clients for a given period of time. When clients report truthfully, reputation information accurately estimates the real QoS delivered by the provider in that period.

At the end of each period, the reputation mechanism publishes the reputation of every provider, and service providers are expected to refund every client the monetary penalty specified in the SLA. The penalty is positive if the delivered QoS (i.e., the reputation) is lower than advertised, and zero otherwise. When the penalty is large enough, [11] proves that rational service providers keep their promises and the market is efficient.

The feedback submitted during one period has no influence over future periods of time. We can therefore separately analyze the reporting incentives of rational clients for each period. Let us consider the time period t where one provider advertises the SLA, $sla_t = \big[p_t, \bar{q}_t, \lambda(\bar{q}_t, R_t)\big]$ where p_t is the price of one service invocation, $\bar{q}_t$ is the advertised QoS, and $\lambda(\bar{q}_t, R_t)$ is the penalty refunded to the client at the end of the period if the reputation R_t is lower than $\bar{q}_t$. For simplicity, we take a linear penalty function:

$$\lambda(\bar{q}_t, R_t) = \max\big(C \cdot p_t \cdot (\bar{q}_t - R_t), 0\big);$$

where C is a constant. We assume that N_t agents require the service from the provider, each of them issuing M service requests during period t. We use the notation a_i, $i = 1, \ldots, N_t$ to refer to one client agent, and a_{-i} to refer to the group of all agents except a_i.

By the end of the period, every agent will have M binary feedback reports, each describing the quality of one service request. Let $o_i = \{s_1, s_2, \ldots, s_M\}, s_j \in \{0, 1\}$ be the feedback set *observed* by agent a_i. As the order of the reports is not relevant, the feedback set can be represented by one integer value counting the number of positive reports: $o_i = \sum_{j=1}^{M} s_j$.

Client agents do not necessarily report true feedback to the reputation mechanism. We denote by $r_i \in \{0, 1, \ldots, M\}$ the feedback set *reported* by agent a_i. Similarly, $r_{-i} \in \{0, 1, \ldots (N_t - 1)M\}$ denotes the union of the feedback sets of the agents a_{-i}.

3.2 Computing the CHI-Scores

At the end of the time period, the reputation mechanism obtains N feedback sets, each containing M binary reports. Based on this information, the reputation mechanism will (a) compute the reputation of the provider for the past period, and (b) reward agents for their feedback such that it is in their best interest to report the truth. For both task, the reputation mechanism will use the CHI-Scores of the feedback sets.

The CHI-Score of the feedback set r_i belonging to the agent a_i is computed based on a Chi-square independence test applied to the sets r_i and r_{-i}. The Chi-square independence test is a well established statistical tool that helps decide if the results of an inquiry are independent from the sample group. Consider, for example, the vacation preferences (seaside or mountains) of the citizens of Switzerland and Italy. m_S out of the N_S inquired Swiss citizens prefers a vacation in the mountains, compared to m_I out of the N_I inquired Italian citizens. Applied to this data, a Chi-square independence test allows us to decide if the vacation preferences of Swiss citizens are really different from those of the Italian citizens.

We use the Chi-square independence test in a similar manner: for every agent, we try to establish if her feedback set is really different from the one of the rest of the agents. However, instead of outputting a binary decision (i.e., reporting pattern is different or not), we modify the test to output a continuous value: the CHI-Score of the feedback set.

We define, χ^i_{Sc}, the CHI-Score of the feedback set r_i as:

$$\begin{aligned} \chi_i^{Sc}(r_i) &= \chi_i^{Sc} = 1 - \mathit{chi2cdf}(\chi_i^2, 1); \\ \chi_i^2 &= \frac{\big(r_i - r_{-i}/(N_t-1)\big)^2}{r_{-i}/(N_t-1)} \\ &\quad + \frac{\big(M - r_i - M + r_{-i}/(N_t-1)\big)^2}{M - r_{-i}/(N_t-1)} \\ &= \frac{M\big(r_i(N_t-1) - r_{-i}\big)^2}{r_{-i}\big(M(N_t-1) - r_{-i}\big)} \end{aligned} \tag{1}$$

where χ_i^2 is the Chi-square value associated with the Chi-square independence test applied to the sets r_i and r_{-i}, and $\mathit{chi2cdf}(\cdot, 1)$ is the chi-square cumulative distribution function with one degree of freedom.

Intuitively, the CHI-Score of r_i is equal to the probability that honest reports of agent a_i could be even further away from r_{-i} than r_i is. The CHI-Score takes values between 0 and 1; the higher the CHI-Score, the "closer" are the sets r_i and r_{-i} and therefore the lower the probability that a_i's reports are distributed differently than the reports of the rest of the agents. We take the CHI-Score to be a direct measure of the agent's honesty, and therefore use it to compute the agent's side-payment, and the weight of the feedback set when computing the reputation of the provider.

3.3 Incentives for Honest Reporting

The reputation of the provider is computed as a weighted average of the feedback sets submitted by the agents:

$$R_t = \frac{1}{M} \frac{\sum_{j=1}^{N_t} \chi_j^{Sc} r_j}{\sum_{j=1}^{N_t} \chi_j^{Sc}}; \tag{2}$$

PROPOSITION 1. *When agents honestly report feedback, the reputation R_t of the provider accurately approximates the actual QoS delivered in the time period t.*

SKETCH OF PROOF. Let q be the real QoS delivered by the service provider during time period t. For all agents a_i, the number of positive reports in the observed feedback set o_i will be distributed according to the binomial distribution with parameters M and q. As all agents report truthfully, the number of positive reports in the set r_{-i} follows a binomial distribution with parameters $(N_t - 1)M$ and q. The expected CHI-Score of agent a_i can be computed from the distributions of r_i and r_{-i}, and is equal for all agents. Since all weights in Equation (2) have the same values, the expected value of the reputation R_t, is equal to the expected ratio of positive reports submitted by all agents: i.e., the delivered QoS, q. ■

Proposition 1 proves that the use of CHI-Scores to weigh feedback sets does not bias the value of the reputation information (i.e., reputation accurately estimates delivered QoS). Therefore, the efficiency results of [11] still apply.

The final price agents pay for the M service invocations in period t is:

$$\begin{aligned} p &= M \cdot \big(p_t - \lambda(\bar{q}_t, R_t)\big) \\ &= M \cdot p_t \Big(1 - max\big(C(\bar{q}_t - R_t), 0\big)\Big); \end{aligned} \tag{3}$$

The price directly depends on the reputation of the provider. By reporting only negative feedback, rational clients decrease the reputation of the provider as much as possible, and thus minimize their payments.

Side-payments for feedback reports can make truth-telling rational. The payment proposed by our mechanism is directly proportional to the CHI-Score of the feedback set submitted by an agent. When all other agents report the truth (hence r_{-i} contains true feedback), a_i expects that r_{-i} follows a similar distribution as o_i, her direct observations. Therefore, the expected CHI-Score of r_i is maximized when $r_i = o_i$: any misreporting increases the "distance" between r_i and r_{-i}. The expected reputation side-payment is therefore maximized when a_i reports honestly; this makes truthful reporting a Nash equilibrium.

The external reward obtained from lying (i.e., the price decrease) imposes supplementary constraints on the reputation side-payments. Truthful reporting remains a Nash Equilibrium only when the price decrease obtained by lying is smaller than the expected side-payment loss. Let $SP(o_i, r_i)$ be the expected side-payment obtained by agent a_i when reporting r_i instead of the observed feedback o_i. $\Delta SP(o_i, r_i) = SP(o_i, o_i) - SP(o_i, r_i)$, $r_i \neq o_i$ is the expected side-payment loss due to misreporting. If $\Delta p(o_i, r_i)$ is the corresponding price decrease, we must have:

$$\Delta SP(o_i, r_i) \geq \Delta p(o_i, r_i); \forall r_i \neq o_i; \tag{4}$$

To compute the potential price decrease obtained by lying, let us consider the effect of misreporting on the reputation of the provider. By reporting r_i instead of o_i, the reputation of the provider becomes:

$$R_t^* = \frac{1}{M} \frac{\chi_i^{Sc}(r_i) r_i + \sum_{j \neq i} \chi_j^{Sc} r_j}{\chi_i^{Sc}(r_i) + \sum_{j \neq i} \chi_j^{Sc}}$$

When N_t is large enough, the misreporting of a_i does not affect the CHI-Scores of other agents, and therefore the expected decrease in reputation can be approximated by:

$$R_t - R_t^* \simeq \frac{o_i \chi_i^{Sc}(o_i) - r_i \chi_i^{Sc}(r_i)}{M \cdot N_t \cdot \chi_i^{Sc}(o_i)}$$

Assuming providers advertise the intended QoS, (i.e., $\bar{q}_t \simeq R_t$), the expected price decrease (due to misreporting) for the M service invocations is:

$$\Delta p(o_i, r_i) = C \cdot p_t \frac{o_i \chi_i^{Sc}(o_i) - r_i \chi_i^{Sc}(r_i)}{N_t \cdot \chi_i^{Sc}(o_i)}$$

Computing the expected side-payment loss due to misreporting requires an explicit model for the beliefs of the agents. We assume agents do not have any prior information about the service provider. Having observed o_i successful service invocations during one time period, the best estimate a_i has about the real QoS delivered by the provider is o_i/M.[5]

We denote by $\phi(r_i, r_{-i})$ the actual payment received by a_i when she reports the feedback set r_i, and the rest of the agents report r_{-i}. Let $\phi(r_i, r_{-i}) = K \cdot \chi_i^{Sc}(r_i)$, where K is a constant, greater than 0. The expected payment of a_i can now be computed by considering the probability distribution of r_{-i} given the real observation o_i:

$$SP(o_i, r_i) = \sum_{r_{-i}=0}^{(N_t-1)M} Pr[r_{-i}|o_i] \phi(r_i, r_{-i}); \tag{5}$$

where $Pr[r_{-i}|o_i]$ follows a binomial distribution for $(N_t - 1)M$ trials with probability o_i/M.

Proposition 2. *Truthful reporting is a Nash Equilibrium when:*

$$K \geq \max_{r_i \neq o_i} \frac{\Delta p(o_i, r_i)}{\Delta SP(o_i, r_i | K = 1)};$$

where $\Delta SP(o_i, r_i|K = 1)$ is the expected side-payment loss for reporting r_i instead of o_i, computed for $K = 1$.

Sketch of Proof. The proof has two parts. First, we show that $\Delta SP(o_i, r_i) \geq 0$ for all $o_i \neq r_i$, and $K > 0$. Then, given the lower bound for K, we show that the price decrease obtained by misreporting is always smaller than the loss caused by the side-payments.

Note that $SP(o_i, r_i) = K \cdot SP(o_i, r_i|K = 1)$ where:

$$SP(o_i, r_i|K = 1) = \sum_{r_{-i}=0}^{(N_t-1)M} Pr[r_{-i}|o_i] \chi_i^{Sc}(r_i, r_{-i});$$

is the expected CHI-Score of the report r_i given the observation o_i. $SP(o_i, r_i|K = 1)$ can be seen as the inter-correlation of two signals: $Pr[x|o_i]$ and $\chi_i^{Sc}(x|r_i)$. When all other agents report the truth, both signals have only one maximum (for $x = o_i(N_t - 1)$ and $x = r_i(N_t - 1)$ respectively) and are (approximately) symmetrically decreasing around the maximum point. The maximum correlation is reached when the maximum points of the two signals are aligned, i.e., $o_i = r_i$, and therefore $SP(o_i, o_i|K = 1) \geq SP(o_i, r_i|K = 1)$ for all $o_i \neq r_i$.

When K satisfies the hypothesis, we have:

$$\begin{aligned} \Delta SP(o_i, r_i) &= K\Big(\Delta SP(o_i, r_i|K = 1)\Big) \\ &\geq \Delta p(o_i, r_i) \text{ for all } o_i \neq r_i; \end{aligned}$$

Therefore it is optimal for a_i to report honestly, given that all other agents also report the truth. This makes truthful reporting a Nash Equilibrium. ■

3.4 Extension to N-ary Feedback

Binary feedback about a transaction can sometime cause the loss of valuable information. For agents using the stock-quote web service, late information (i.e., received after the

[5] for $o_i = 0$ and $o_i = M$, the actual estimate of the QoS is actually ε respectively $1 - \varepsilon$ for ε small

deadline) could still be more valuable than no information at all. One can therefore imagine a finer-grained reporting system where the feedback also gives some information about the delay of the answer: one feedback is submitted when the request was answered on time, another when the answer was no more than 1 minute late, etc.

In general, one can consider an N-ary reporting scheme where feedback can take any value from the predefined set $\mathcal{S} = \{s_1, \ldots s_S\}$ of cardinality S. Agents share a common ontology about the elements of $\mathcal{S}$ (i.e., they interpret feedback in the same way). The feedback set observed by an agent becomes a vector $o_i = (o_i^1, \ldots, o_i^S)$ where o_i^j is the number of service requests where a_i perceived the feedback value s_j. Similarly, the reported feedback set of an agent (i.e., r_i), as well as the feedback set of all other agents (i.e., r_{-i}) are vectors with S elements. The QoS and the reputation of the provider become probability distributions over the possible feedback values: $\bar{q}_t = (q_{s_1}, \ldots, q_{s_S}) \in [0,1]^S$ such that $\sum_{j=1}^{S} q_{s_j} = 1$.

The extension of the mechanism to N-ary feedback is straightforward: the CHI-scores are computed in the same way and become:

$$\chi_i^{Sc}(r_i) = \chi_i^{Sc} = 1 - chi2cdf(\chi_i^2, S-1);$$
$$\chi_i^2 = \sum_{j=1}^{S} \frac{\left(r_i^j(N_t-1) - r_{-i}^j\right)^2}{(N_t-1)r_{-i}^j}$$

where the Chi-square value is a summation over the dimensions of the feedback set, and the chi-square cumulative distribution function has $S-1$ degrees of freedom. The qualitative results of Proposition 1 and Proposition 2 remain the same.

4. EXPERIMENTAL EVALUATION

In this section we experimentally evaluate the mechanism described above, using several criteria. First, we study the accuracy of the reputation information in the presence of irrational clients that misreport feedback regardless of the monetary loss they incur. Second, we investigate the robustness of the honest-reporting incentives in the presence of such irrational agents. Finally, we compare the cost of our mechanism against existing side-payment schemes.

We consider the same example presented in the introduction, where a web service delivers stock quotes to interested agents. Every morning, the web service advertises a SLA specifying the price of a service request, the promised QoS (as a probability of delivering the information within a given deadline), and the coefficient C of the linear reputation-based penalty function λ. Clients subscribe to the service before the opening time of the market. Each subscription specifies (a) the code of the desired share and (b) the set of M time points when an information update is requested. An example of a subscription is: *"Request the price of the share X every 15 minutes starting at 9:00 AM and ending at 5:PM"*. Every subscription generates M service invocations.

The real QoS delivered to clients depends on the infrastructure of the provider and on the number of accepted requests. The provider can choose the intended QoS by limiting the total number of subscribers for a given day. When the stock market opens, all requests have already been registered and the delivered QoS remains fixed for the entire day. The reputation information is computed by the reputation mechanism at the end of every day, and all penalties are paid to the clients before the beginning of the next day.

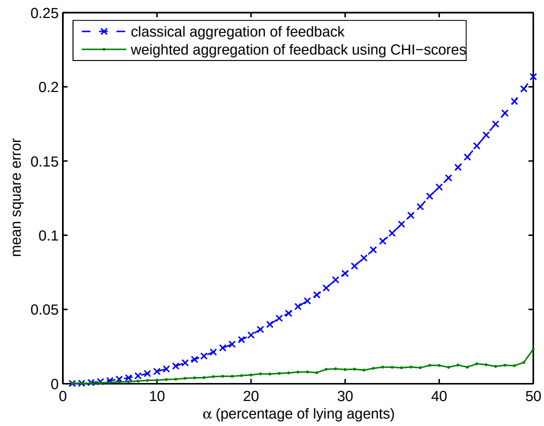

Figure 2: Mean square error of reputation information.

We take one day and one web service. We assume the total number of clients is N, and that all client requests are similar: the price of a share is requested every m minutes. The total number of requests per client during one day is therefore $M = 8hours \times 60min/m$.

As a first experiment, we study the accuracy of the reputation mechanism when irrational clients lie, despite the guarantees provided by Proposition 2. Figure 2 plots the mean square error between the expected and real reputation information when α percent of the agents lie (submit only negative feedback). The performance of our mechanism is compared against a classical (un-weighted) aggregation scheme for $N = 100$ agents, $M = 40$ interactions per agent, and the real QoS equal to 0.9. The average is taken over 100 runs. As it can be seen, using CHI-Scores to weigh feedback sets dramatically improves the robustness of the mechanism against lying agents.

Irrational lying also affects the honest reporting incentives for the rational agents. The most extreme lying strategy (worst-case scenario) is when irrational agents constantly submit negative feedback. Depending on the number of irrational reporters, rational agents find it optimal to slightly down-bias their real report. Figure 3 plots the probability that a rational agent will report something else than the true observation as the number of lying agents increases. For this experiment we took the same delivered quality value (i.e QoS = 0.9) but we considered only $M = 20$ interactions per client. The total number of clients varied from $N = 100$ to $N = 400$. The graph reveals an almost instantaneous switch from the truthful strategy to some lying strategy. For up to 2.5% liars, truth-telling is optimal; from that point onward, every supplementary liar makes down-biasing rational for an increasing set of observations (some observations are still reported truthfully, but the probability of such observations appearing decreases rapidly with the number of liars). With 3.5% liars, down-biassing becomes rational for all observations. This low threshold is mostly due to the extreme lying strategy we have considered. In reality, dishonest agents will lie only once every L reports, and consequently, the toler-

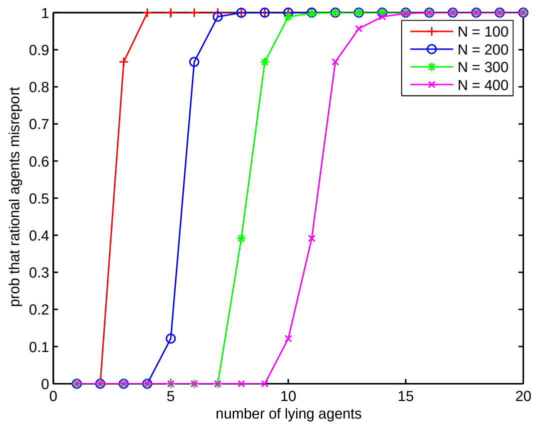

Figure 3: Misreporting incentives in the presence of irrational liars.

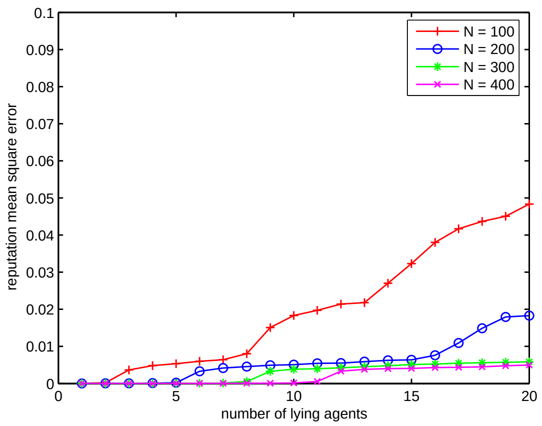

Figure 4: Reputation mean square error when rational agents are aware of the presence of irrational reporters.

ance threshold will be multiplied by L.

Even when truth-telling is not optimal, the submitted reports are not too far from the truth, and thus do not greatly distort reputation information. Figure 4 plots the corresponding mean square error for the reputation value when rational agents are aware of the presence of irrational reporters and consequently adapt their reports.

The reputation side-payments contribute substantially to the global cost of the reputation mechanism. In theory, all side-payment mechanisms can be made budget balanced by asking fixed (or periodically updated) fees from the client agents (i.e., clients also pay for using the reputation mechanism). However, none of the existing online markets charges clients for access to reputation information. The cost of reputation management is absorbed by the market and later distributed to service providers. Moreover, Dellarocas [4] theoretically proves that increasing amounts of truthful feedback helps rational service providers cut down costs (e.g., advertising costs) and thus justifies the providers' contribution to reputation management costs.

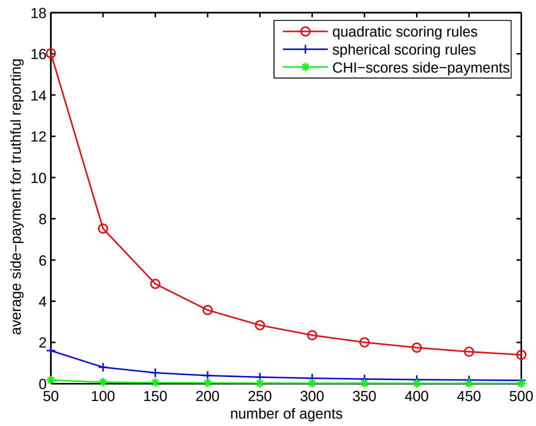

Figure 5: Average side-payment per client per transaction for truthful reporting.

It therefore becomes an important objective for mechanism designers to minimize such costs, by minimizing the side-payments. We compare the costs of our mechanism with existing side-payment schemes based on scoring rules [12, 11]. We consider two of the best-known scoring rules: the spherical and quadratical scoring rule. The resulting side-payments are scaled (by addition and multiplication with a constant) to the minimum positive side-payments that counterbalance the price decrease obtained from misreporting. We take the average expected side-payment for truthful reporting, $\bar{SP} = \sum_{o_i} SP(o_i, o_i)$, per client, per transaction. Figure 5 plots the averages for the three mechanisms when $M = 20$ interactions per day per client, the price of service and the penalty constant are unitary (i.e., $p_t = C = 1$) and N varies from 50 to 500 clients per day. The side-payments based on CHI-Scores are one order of magnitude lower than the cheapest of the side-payments based on the considered scoring rules. When one service invocation costs 1 dollar, the average side-payment based on CHI-Scores is 7.5 cents for $N = 50$ clients and 1.3 cents for $N = 500$ clients. In the same conditions, the average side-payment computed using the spherical scoring rule is 79 respectively 15 cents.

5. DISCUSSION AND FUTURE WORK

One of the important assumptions behind our mechanism is that all customers in a certain group receive the same QoS within one period of time. In real systems, however, request arrival rates as well as other real-time factors might trigger non-negligible short term fluctuations in the service quality. Certain clients will thus receive better service than others.

However, most of these fluctuations will look random for the clients. Even though quality observations are no longer independently identically distributed, one client, due to the limited number of observations, will not be able to model the stochastic behavior of the service provider. Thus, the clients' beliefs remain as described in this paper, and the CHI-Scores encourage truthful reporting.

On the other hand, easily-identifiable QoS fluctuations patterns should not exist in a rational market. When service providers can identify the clients that are likely to receive higher quality, it makes sense to treat those clients as part of a different, "preferred" group. Providers can thus offer different conditions (QoS and price) to each group, and clients within the same group continue to be treated "approximately equally". As a last resort, the mechanism designer can of course solve the problem by shortening the length of the time period.

A related problem comes from subjective interpretations of quality. Different agents might have different criteria for high quality, and this will affect our scoring mechanism. For this reason, our paper targets specific markets where autonomous agents request (web-)services. Interactions in such environments are usually regulated by SLA's that consider objective measures of QoS (such as response time, availability, etc). Dealing with subjective perceptions is a challenging direction for future work.

Truthful reporting is guaranteed as a Nash equilibrium by Proposition 2. However, honest reporting is not the only equilibrium strategy: e.g., reporting only negative experiences is also a Nash equilibrium. The problem of multiple reporting equilibria has been formally identified in [9]. Fortunately, there is a simple way of enforcing truth-telling by using a small number of trusted, third-party reports.

As future work, we plan to extend this mechanism along several dimensions. First, we plan to increase the robustness of the mechanism against irrational liars. In a future version we envisage to rate the feedback set of one agent against the reports of a fraction (not all) of the remaining agents. By carefully choosing the "reference" set for every agent, we plan to increase the threshold for the percentage of tolerated irrational reports.

Second, we plan to extend the set of contexts where the mechanism can be successfully used, by allowing some limited amounts of external information. As before, we expect to be able to group agents according to their prior information, and rate the feedback set of an agent against the reports submitted by other agents in the same group.

Third, we plan to investigate other distance metrics for scoring set of reports.

Last, but not least, we plan to address the problem of collusion. In the present mechanism, two or more agents can benefit by coordinating their false reports. To solve this problem we plan using the idea in [11] where a small number of trusted reports (i.e., reports that are true with high probability) can deter collusion.

6. CONCLUSION

When client agents can submit several feedback reports within a certain period of time, existing payment schemes that reward honest feedback are either inefficient, or extremely expensive. In this paper we present a novel side-payment scheme based on CHI-Scores. We consider the entire feedback set of an agent, not her individual reports. The CHI-Score of a feedback set results from a Chi-square independence test and reflects the correlation between the reports of one agent and those of the rest of the group. The score is used to weigh feedback sets when computing reputation information, and to scale the reputation side-payments. We prove that such monetary rewards make truth-telling rational, and show how external benefits obtained from lying can be counterbalanced. The mechanism can be extended to finer-grained feedback systems, and generates significantly lower costs. Experimental results show that a small number of irrational liars, does not significantly degrade the performance of the mechanism.

7. REFERENCES

[1] R. Cooke. *Experts in Uncertainity: Opinion and Subjective Probability in Science*. Oxford University Press: New York, 1991.

[2] A. Dan, D. Davis, R. Kearney, A. Keller, R. King, D. Kuebler, H. Ludwig, M. Polan, M. Spreitzer, and A. Youseff. Web services on demand: WSLA-driven automated management. *IBM Systems Journal*, 43(1):136–158, 2004.

[3] C. Dellarocas. Sanctioning Reputation Mechanisms in Online Trading Environments with Moral Hazard. MIT Sloan Working Paper #4297-03, 2004.

[4] C. Dellarocas. Strategic Manipulation of Internet Opinion Forums: Implications for Consumers and Firms. Working Paper, 2005.

[5] D. Houser and J. Wooders. Reputation in Auctions: Theory and Evidence from eBay. *Journal of Economics and Management Strategy*, forthcoming, 2005.

[6] R. Ismail and A. Jøsang. The Beta Reputation System. In *Proceedings of the 15th Bled Conf. on E-Commerce*, 2002.

[7] R. Jurca and B. Faltings. An Incentive-Compatible Reputation Mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.

[8] R. Jurca and B. Faltings. "CONFESS". Eliciting Honest Feedback without Independent Verification Authorities. In *Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI 2004)*, New York, USA, July 19 2004.

[9] R. Jurca and B. Faltings. Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms. In *Internet and Network Economics*, volume 3828 of *LNCS*, pages 268 – 277. 2005.

[10] R. Jurca and B. Faltings. Reputation-based Pricing of P2P Services. In *Proceedings of the Wokshop on Economics of P2P Systems*, Philadelphia, USA, 2005.

[11] R. Jurca and B. Faltings. Reputation-based Service Level Agreements for Web Services. In *Service Oriented Computing (ICSOC - 2005)*, volume 3826 of *LNCS*, pages 396 – 409. 2005.

[12] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. Forthcoming in Management Science, 2005.

[13] T. G. Papaioannou and G. D. Stamoulis. An Incentives' Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems. In *Proceedings of IEEE/ACM CCGRID 2005*, 2005.

[14] D. Prelec. A bayesian truth serum for subjective data. *Science*, 306(5695):462–466, 2004.

[15] P. Resnick, R. Zeckhauser, J. Swanson, and K. Lockwood. The Value of Reputation on eBay: A Controlled Experiment. *Experimental Economics*, forthcoming, 2005.

[16] L. Teacy, J. Patel, N. Jennings, and M. Luck. Coping with Inaccurate Reputation Sources: Experimental Analysis of a Probabilistic Trust Model. In *Proceedings of AAMAS*, Utrecht, The Netherlands, 2005.

[17] K. Walsh and E. Sirer. Fighting Peer-to-Peer SPAM and Decoys with Object Reputation. In *Proceedings of P2PECON*, Philadelphia, USA, 2005.

[18] N. A. Weiss. *Elementary Statistics*. Addison-Wesley Publishing Company, 1996.

[19] A. Whitby, A. Jøsang, and J. Indulska. Filtering out Unfair Ratings in Bayesian Reputation Systems. In *Proceedings of the 7th Intl. Workshop on Trust in Agent Societies*, 2004.

[20] B. Yu and M. Singh. Detecting Deception in Reputation Management. In *Proceedings of the AAMAS*, Melbourne, Australia, 2003.