

Enforcing Truthful Strategies in Incentive Compatible Reputation Mechanisms

Radu Jurca and Boi Faltings

Ecole Polytechnique Fédérale de Lausanne (EPFL)
Artificial Intelligence Laboratory
CH-1015 Lausanne, Switzerland
{radu.jurca, boi.faltings}@epfl.ch

Abstract. We commonly use the experience of others when taking decisions. Reputation mechanisms aggregate in a formal way the feedback collected from peers and compute the *reputation* of products, services, or providers. The success of reputation mechanisms is however conditioned on obtaining true feedback. Side-payments (i.e. agents get paid for submitting feedback) can make honest reporting rational (i.e. Nash equilibrium). Unfortunately, known schemes also have other Nash equilibria that imply lying. In this paper we analyze the equilibria of two incentive-compatible reputation mechanisms and investigate how undesired equilibrium points can be eliminated by using trusted reports.

1 Introduction

In a world that offers an ever increasing number of choices, we commonly use the experience of our peers when making decisions. The feedback coming from previous users can be aggregated into the *reputation* of a product, service or manufacturer, and accounts for the data that cannot be directly observed before the purchase: e.g. reliability, technical support, etc.

As reputation mechanisms become more and more popular in online markets, it is important to ensure that selfish agents have the right incentives to report honest feedback (i.e. the mechanism is incentive compatible). One way to elicit truthful information is to pay the reports according to their estimated truthfulness. Since objective verification is usually impossible, the truthfulness of a report is assessed by comparing it with other reports coming from peers. When the observations of different clients are sufficiently correlated, there exist payment rules that make truthful reporting be Nash Equilibrium (NEQ): i.e. rational agents report the truth given that all other agents report the truth. [9] and [6] describe concrete mechanisms.

Unfortunately, such mechanisms also have other NEQ points where agents lie. The existence of multiple equilibria is a serious problem when engineering real reputation mechanisms: nothing can guarantee that the desired (i.e. truthful) equilibrium strategy is selected. Moreover, the payoff generated by the truthful equilibrium is often dominated by the payoff in a non-truthful equilibrium. Thus, the selection of the truthful strategy becomes even more problematic.

In this paper we propose a method of enforcing the selection of the truthful strategy based on trusted reports (i.e. verifiable reports coming from specialized reporters). Such reports can constitute a true reference base against which other feedback can be evaluated. When enough trusted reports are available, the incentive compatible NEQ becomes unique.

Efficiency, however, dictates that the number of trusted reports required be kept as small as possible. We therefore investigate the reputation mechanisms described in [9] and [6], and derive analytical and numerical solutions for the minimum percentage of trusted reports required to enforce the truthful strategies. Besides the results for the two specific mechanisms, the paper introduces a general methodology for eliminating undesired equilibrium points, and offers insights into the dynamics of feedback reporting mechanisms. Sections 2 and 3 briefly introduce the two reputation mechanisms. Section 4 analyzes the set of Nash equilibrium points and analytically shows how trusted reports can be used to eliminate the undesired equilibria. Numerical results are presented and interpreted in Section 5, followed by related work and a conclusion.

2 The MRZ Incentive Compatible Reputation Mechanism

In [9], Miller, Resnick and Zeckhauser (henceforth referred to as MRZ) consider that a number of clients sequentially experience the same product whose *type*¹ is drawn from a set of possible types T .²

The real type of the product is unknown to clients and does not change during the experiment. After every interaction, the client observes one signal s (from a set of possible signals, S , of cardinality M) about the type of the product. The observed signals are independently identically distributed conditioned on the real type t of the product. $f(s_i|t)$ denotes the probability that the signal s_i is observed when the product is of type t . $\sum_{s_i \in S} f(s_i|t) = 1$ for all $t \in T$.

After every interaction, the client is asked to submit feedback about the signal she has observed. A reputation mechanism collects the reports, and updates the probability distribution over the possible types of the product. Let p characterize the current belief of the reputation mechanism (and therefore of all agents that can access the reputation information) about the probability distribution over types. $p(t)$ is the probability that the product is of type t , and $\sum_{t \in T} p(t) = 1$. When the reputation mechanism receives a report $r \in S$, the posterior belief is updated using Bayes' Law:

$$p(t|r) = \frac{f(r|t) \cdot p(t)}{Pr[r]};$$

¹ The type of a product defines the totality of relevant characteristics of that product. e.g. quality, reliability, etc.

² The set of possible types is the combination of all values of the attributes that define the type. While this definition generates an infinite-size set of types, in most practical situations, approximations make the set of possible types countable. For example, the set of possible types could have only two elements: *good* and *bad*.

where $Pr[r] = \sum_{t \in T} f(r|t) \cdot p(t)$ is the probability of observing the signal r .

Every feedback is paid according to a payment rule that takes into account the current belief, the value of the report and the value of another future report submitted by some other client (called the *rater*). When this payment is defined by a *proper scoring rule*³, MRZ show that every agent has the incentive to submit the true feedback given that the rater also reports honestly. The mechanism thus has an incentive-compatible NEQ.

One possible payment rule is:

$$R(p(\cdot), r, r_r) = \log(Pr[r_r|r, p(\cdot)]) = \log\left(\sum_{t \in T} p(t|r) \cdot f(r_r|t)\right);$$

where $p(\cdot)$ is the prior belief of the agent (and of the reputation mechanism), $r \in S$ is the report of the agent, $r_r \in S$ is the future report of the designated rater, and $Pr[r_r|r]$ is the posterior probability that the signal r_r will be observed by the rater given that r was observed by the current reporter. When denoting the payment received by an agent, we will frequently ignore the dependence on the belief and have $R(s_i, s_j)$ represent the payment received by an agent reporting s_i when the rater reports s_j .

3 The JF Incentive-Compatible Reputation Mechanism

Jurca and Faltings (henceforth referred to as JF) describe in [6] an incentive compatible reputation mechanism in a setting where the binary signal observed by the clients is influenced not only by the type of the service, but also by time. The probability distribution of the observed signal is thus modeled by a Markov chain of variable length.

The side-payment for reports follows a very simple rule, and does not depend on the beliefs of the agent or those of the reputation mechanism. A report is paid only if the next report submitted by some other client about the same service has the same value. The amount of the payment is dynamically scaled such that the mechanism is budget-balanced.

The Markov model for the observable signals is very appropriate for services offered by software agents, where failures are correlated. If we take the example from the previous section, and consider that the product is actually a service provided by a webservice, it is very unlikely that individual signals follow an independent distribution. A service failure due to a software or hardware problem is likely to attract other service failures in the immediate future. Likewise, a present successful invocation signals that everything works well with the infrastructure, and is probably going to be followed by other successful service invocations.

While the MRZ mechanism can also be adapted for Markov models of behavior, it requires that the model be common knowledge among the agents: i.e. all agents must agree on the length of the model and on the fact that there is

³ see [4] for an introduction to the scoring rules

a unique set of parameters characterizing that model. MRZ argue that private information can be accommodated by requiring the agents to first submit their private information, and then report the feedback. The computation of the payment will take into account the reported private information, and will make it rational for the agents to truthfully submit feedback afterwards.

However, reporting private information and feedback introduces additional cheating opportunities. Although no agent can obtain an expected payoff greater than the one rewarded by the truthful strategy, malicious reporters can bias the declared private information in order to make any desired report (weakly) optimal: e.g. an agent willing to bad-mouth a provider, can do so without being penalized by submitting appropriately modified private information.

Having side-payments that do not depend on the beliefs of the agents, the JF mechanism allows the agents to have private beliefs about the model of the webservice, as long as these beliefs satisfy some general constraints. Of course, the freedom of having private beliefs is paid by the constraints that limit the contexts in which incentive-compatibility is guaranteed.

4 Equilibrium Strategies

Formally, a reporting strategy of an agent a is a mapping σ from the set of signals S to the set ΔS containing all probabilistic combinations of signals from S . $\sigma(s_i) = \sum_{j \in S} \alpha_j^i s_j$ denotes that an agent a following reporting strategy σ , will report s_j with probability α_j^i given that the signal observed was s_i . $\sum_{j=1}^M \alpha_j^i = 1$ for all $i \in \{1, \dots, M\}$. The set of all reporting strategies is denoted as \mathcal{S} .

Let σ^* be the incentive-compatible strategy, i.e. $\sigma^*(s_i) = s_i$ for all $s_i \in S$. By an abuse of notation we also use s_j to denote the ‘‘constant’’ reporting strategy (i.e. $s_j(s_i) = s_j$ for all $s_i \in S$) and ΔS to denote the set of all reporting strategies.

When a uses reporting strategy σ and her rater (i.e. a_r) uses the reporting strategy $\sigma' = (\beta_j^i)$, the expected payment of a when observing the signal s_i is:

$$E[\sigma, \sigma', s_i] = \sum_{j=1}^M \alpha_j^i \left(\sum_{k=1}^M Pr[s_k | s_i] \cdot \left(\sum_{l=1}^M \beta_l^k \cdot R(s_j, s_l) \right) \right); \quad (1)$$

where $Pr[s_k | s_i]$ is the probability that the rater observes the signal s_k given that a has observed s_i , and the function $R(s_j, s_l)$ gives the payment made by the reputation mechanism to a when a reports the signal s_j and a_r reports s_l . For the MRZ mechanism the function $R(s_j, s_l)$ is given by one scoring rule. For the JF mechanism, the function $R(s_j, s_l)$ is 1 if $s_l = s_j$ and 0 otherwise.

Definition 1. *A reporting strategy σ is a NEQ of the reputation mechanism iff $\forall s_i \in S$, no agent deviates from σ , as long as her rater reports according to σ . Formally, σ satisfies: $E[\sigma, \sigma, s_i] \geq E[\sigma', \sigma, s_i], \forall s_i \in S, \sigma' \neq \sigma$.*

Definition 1 restricts possible reporting strategies to *symmetric* ones. General N -player feedback reporting games might have asymmetric Nash equilibria as well (i.e. every agent uses a different reporting strategy). However, online

markets usually assume an infinite number of anonymous clients providing feedback. In this case, all reporting equilibria are symmetric. To prove that, assume an asymmetric equilibrium, and two agents, a_i and a_j using different reporting strategies: $\sigma_i \neq \sigma_j$. For both a_i and a_j the rater will be drawn from the same (infinite) set of future reporters. Therefore, for both a_i and a_j , the rater's strategy will be the same strategy σ , computed as a mix of the strategies of all future reporters. Then, $E[\sigma_i, \sigma, \cdot] \geq E[\sigma_j, \sigma, \cdot]$ as σ_i is optimal for a_i and $E[\sigma_j, \sigma, \cdot] \geq E[\sigma_i, \sigma, \cdot]$ as σ_j is optimal for a_j . Consequently, $\sigma_i = \sigma_j$, and by induction all clients use the same strategy in equilibrium (i.e. the equilibrium reporting strategy is symmetric).

Both the MRZ and the JF mechanisms have many NEQ strategies. In general, finding all NEQ points of a game is a difficult problem [3]. However, for the special case of binary reputation mechanisms, Proposition 1 completely characterizes the set of equilibrium strategies.

Proposition 1. *Given a binary incentive-compatible reputation mechanism, always reporting positive feedback and always reporting negative feedback are Nash equilibria. At least one of these equilibria generates a higher payoff than the truthful equilibrium.*

Proof. Let “+” and “-” denote the positive and respectively the negative quality signals. The mechanism is incentive-compatible, so $E[\sigma^*, \sigma^*, +] \geq E[-, \sigma^*, +]$ and $E[\sigma^*, \sigma^*, -] \geq E[-, \sigma^*, -]$. Expanding $E[\cdot, \cdot, \cdot]$ and taking into account that $Pr[+|+] \geq Pr[+|-]$ (easy verifiable by applying Bayes law) we obtain: $R(+, +) \geq R(-, +)$ and $R(-, -) \geq R(+, -)$. Therefore, $E[+, +, \cdot] \geq E[-, +, \cdot]$ and $E[-, -, \cdot] \geq E[+, -, \cdot]$, and thus, according to Definition 1, the strategies + and - (i.e. always reporting positive, respectively negative feedback) are NEQ.

Let $\rho = \max(R(+, +), R(-, -))$. Then,

$$\begin{aligned} E[\sigma^*, \sigma^*, +] &= Pr[+|+]R(+, +) + Pr[-|+]R(+, -) \\ &\leq Pr[+|+]R(+, +) + Pr[-|+]R(-, -) \leq \rho; \end{aligned}$$

Similarly, $E[\sigma^*, \sigma^*, -] \leq \rho$, therefore, at least one of the constant reporting NEQ strategies generates a higher expected payoff than the truthful equilibrium. \square

The results of Proposition 1 are valid for all binary incentive-compatible reputation mechanisms, and prove that honesty is always dominated by at least one of the constant reporting strategies. We conjecture the existence of a similar result for all IC reputation mechanisms.

4.1 The influence of trusted reports

For all incentive compatible reputation mechanisms, the truthful reporting strategy σ^* , is a strict Nash equilibrium. When the report submitted by the rater is always trusted (i.e. true), the expected payment received by a , given that she has observed the signal s_i and uses the reporting strategy σ , is $E[\sigma, \sigma^*, s_i]$. As the

rater's strategy is fixed, the only Nash equilibrium strategy of a is the truthful reporting strategy σ^* . Any other reporting strategy will generate a strictly lower payoff (Definition 1).

Since trusted reports are expensive, it is interesting to see if undesired Nash equilibrium points can be eliminated by using only a *probabilistic* rating against a trusted report: i.e. a report is rated with probability q against a trusted report and with probability $1 - q$ against a normal report. The expected payoff to a from the equilibrium strategy σ , given that she has observed the signal s_i is then:

$$E_q[\sigma, \sigma, s_i] = q \cdot E[\sigma, \sigma^*, s_i] + (1 - q) \cdot E[\sigma, \sigma, s_i]$$

The strategy σ continues to be a Nash equilibrium strategy if and only if for all other reporting strategies σ' , $E_q[\sigma', \sigma, s_i] < E_q[\sigma, \sigma, s_i]$, for all signals s_i . Finding the minimum probability q such that the incentive-compatible reporting strategy remains the only Nash equilibrium strategy of the mechanism involves solving the following problem:

Problem 1. Find the minimum $q^* \in [0, 1]$ such that for all q , $q^* \leq q \leq 1$, for all reporting strategies $\sigma \neq \sigma^*$, there is a signal s_i and a strategy $\sigma' \neq \sigma$ such that $E_q[\sigma, \sigma, s_i] < E_q[\sigma', \sigma, s_i]$.

Problem 1 is hard to solve in the general case, and its result are very restrictive. A relaxation would be to eliminate only those equilibrium strategies that generate a higher payoff than the incentive compatible strategy. The practical justification for this relaxation is that rational agents always choose from a set of possible equilibrium strategies the one that generates the highest payoff. Given that truthful reporting yields the highest payoff, we argue that it is not necessary from a practical perspective to eliminate all other Nash equilibrium points.

Finding the minimum probability, q^* , such that the incentive-compatible reporting strategy generates the highest payoff implies solving the following problem:

Problem 2. Find $q^* = \min(q)$, s.t. $f(q^*) = 0$, where $f(q) = \max_{\sigma, s_i} E_q[\sigma, \sigma, s_i] - E[\sigma^*, \sigma^*, s_i]$ s.t. σ is a NEQ: i.e. $E_q[\sigma, \sigma, s_k] \geq E_q[s_j, \sigma, s_k]$ for all $s_j, s_k \in S$.

Problem 2 contains two nested optimizations: (1) finding the Nash equilibrium strategy that generates the highest payoff, and (2) finding the minimum value of q (i.e. q^*) for which the highest Nash equilibrium payoff corresponds to the incentive-compatible reporting strategy. Finding the highest Nash equilibrium payoff is a NP-hard problem [3]. The function $f(q)$, on the other hand, is decreasing in q and therefore a binary search can be used to find the minimum value of q . Please note that the solutions to problem 2 also represent lower bounds for the solutions of problem 1.

The solution q^* of Problem 2 does not necessarily represent the overall percentage of trusted reports needed by the reputation mechanism. For example, the MRZ mechanism allows to reuse trusted reports. The same trusted report can be used to assess the honesty of more than one feedback. In extremis, one

could imagine that the same trusted report is used to assess all other reports collected by the reputation mechanism. The actual percentage of trusted reports needed by the mechanism is hence very low, and equal to the value of q^* divided by the total number of reports that are rated against the same trusted report.

Using the same trusted report poses, however, some problems. First, all feedback has to be rated in the same time, i.e. after all reports have been submitted. This delays the side-payments to clients, and weakens the monetary incentives to report. Second, the trusted report can become outdated. In a dynamic system (service providers change their type by updating for example their infrastructure), the validity of any report is limited to a certain time window. Third, it leaves the mechanism vulnerable to mistakes of the trusted reporters.

In practice, a periodically updated set of trusted reports can be used to rate normal feedback. For any report, one trusted report can be randomly chosen from this set. Thus, a compromise can be reached between cost and stronger incentives to report the truth.

The JF mechanism, on the other hand, requires a fresh rater for every submitted report (i.e. the report of the next agent is always used to rate the present feedback). Therefore, the threshold value q^* that a report is rated against a trusted report becomes the overall percentage of trusted reports needed by the reputation mechanism.

Our analysis also offers an interesting insight for influencing the behavior of incentive compatible reputation mechanisms. From a dynamic perspective, the reputation mechanism can shift from one Nash reporting equilibrium to another. A mechanism operator will therefore be interested to know how easy it is to switch the reporting strategy from lying to truth-telling.

For example, let us take a reputation mechanism currently coordinated on an equilibrium σ . Assuming that the operator can observe σ , he can shift the reporting equilibrium to the truthful one by publicly committing to rate every report against a trusted report with probability $\bar{q} = \max_{s_i} \frac{E[\sigma, \sigma, s_i] - E[\sigma^*, \sigma^*, s_i]}{E[\sigma, \sigma, s_i] - E[\sigma, \sigma^*, s_i]}$. Note that $\bar{q} \leq q^*$, with equality holding when σ is the reporting NEQ yielding the highest payoff. Thus, in particular, the operator can drive the reputation mechanism to the truthful equilibrium in the beginning.

Once the agents have coordinated on the truthful strategy (by some external coercion or natural initiative), the operator can stop using trusted reports. It will take a significant proportion of deviators (i.e. at least $1 - q^*$) coordinated on a different payoff equilibrium in order to make it rational for the other agents to also switch to this non-truthful equilibrium. Since q^* is often quite low, this can be very useful in practice. This observation opens new research and design opportunities for practical reputation mechanisms where external intervention (e.g. trusted reports) is only periodically needed in order to (re)coordinate a sufficient proportion of agents on the desired strategy.

5 Numerical Analysis

For the JF mechanism the incentive-compatible equilibrium dominates all other equilibria only when feedback is rated against a trusted report with probability

greater than $q^* = \max\left(\frac{1-Pr[-|-]}{Pr[-|-]}, \frac{1-Pr[+|+]}{Pr[+|+]}\right)$, where $Pr[+|+]$ and $Pr[-|-]$ are the probabilities of observing a positive (respectively negative) signal in the next round given a positive (respectively negative) signal observed in the present. Both $Pr[-|-]$ and $Pr[+|+]$ vary in the interval $[0.5, 1]$ (smaller values for $Pr[-|-]$ or $Pr[+|+]$ are not allowed by the assumptions of the JF mechanism), therefore q^* can take any value between 0 and 1. For example, a webservice modeled by a Markov chain with transition probabilities $Pr[+|+] = 0.95$ and $Pr[-|-] = 0.9^4$ requires $q^* = 1/9$. The greater the correlation between the behavior of the webservice in successive transactions, the lower the threshold value, q^* . In such cases, there is little uncertainty about the signal observed by the rater, and therefore the incentive-compatible strategy of the JF mechanism yields payoffs that approach the maximum possible payoff.

For the MRZ mechanism we numerically investigate settings with N possible product types and N observable signals. Each signal characterizes one type, and uniform noise “scrambles” the observation of clients. The conditional probability distribution of signals is defined as $f(s_j|t_i) = 1 - \delta$ when $s_j = s_i$ and $f(s_j|t_i) = \frac{\delta}{N-1}$ when $s_j \neq s_i$. $\delta = 10\%$ is the “level” of noise.

For $N = 2$, Figure 1 plots the threshold value q^* against the set of possible beliefs of the clients (characterized by the prior probability assigned to the *good* type). The values of q^* range between 0.6 and 0.8. The gaps at both ends of the interval are explained by the “activation” of previously inefficient reporting strategies. When one type is very probable (e.g. the good type), the constant reporting strategy “-” is very inefficient. $Pr[-|-]$ is very small, therefore the payoff generated by always reporting negative feedback is lower than the payoff rewarded to the truthful strategy. As soon as the prior probability of the *bad* type becomes big enough, both constant reporting strategies are more profitable than the truthful strategy, hence more trusted reports are needed to enforce truth-telling.

For $N = 3$ the space of possible beliefs is two dimensional: two probabilities entirely characterizes the prior distribution over the three types. Figure 2 presents three slices through the 3-dimensional graph for three different probabilities of the type t_1 : $p(t_1) = 0.1$, $p(t_1) = 0.3$ and $p(t_1) = 0.5$. The threshold value q^* varies between 0.5 and 0.8.

For higher number of types, the graphical representation of the space of beliefs becomes impossible. Moreover, solving the optimization problem for an increasing number of types (and therefore signals) becomes exponentially more difficult. However, the solution for $N = 4, 5, 6$ types and beliefs normally distributed around every type does not bring any surprises. As in the previous cases, the values of q^* range from 0.5 to 0.8, with higher values for more focused beliefs, and lower values for increased ambiguity in beliefs.

⁴ $Pr[+|+]$ is the probability of successful service at time $t+1$ given a successful service at time t ; similarly $Pr[-|-]$ is the probability of service failure at time $t+1$ given a failure at time t




Fig. 1. Threshold value q^* for the MRZ mechanism, when $N = 2$ and $\delta = 10\%$.




Fig. 2. Threshold value q^* for the MRZ mechanism, when $N = 3$, and $\delta = 10\%$.

6 Related work

In their seminal papers, Kreps, Wilson, Milgrom and Roberts [8] prove that cooperative equilibria can exist in finitely repeated games due to the reputation effect. Since then, numerous computational reputation mechanisms have been described, ranging from mechanisms based on direct interactions [2] to complex social networks [11] where agents ask and give recommendations to their peers. Centralized implementations as well as completely decentralized [1] have been investigated.

Besides the two mechanisms treated in this paper, a number of other mechanisms address the problem of eliciting honest feedback from self interested participants. For e-Bay-like auctions, the Goodwill Hunting mechanism [5] provides a way to make the sellers indifferent between lying or truthfully declaring the quality of the good offered for sale. Momentary gains or losses obtained from misrepresenting the good's quality are later compensated by the mechanism which has the power to modify the announcement of the seller.

Jurca and Faltings [7] take a different approach and achieve in equilibrium truthful reporting by comparing the two reports coming from the buyer and the seller involved in the same transaction. Using the same idea, Papaioannou and Stamoulis [10] describe a mechanism suitable for P2P environments.

This paper also relates to the vast literature concerned with computing Nash equilibrium strategies [3] and with the ongoing efforts of the networking community to design routing algorithms that have a unique Nash equilibrium point with the desired properties.

7 Conclusion

Obtaining true feedback is of vital importance to the success of online reputation mechanisms. When objective verification is not available, economic measures must exist to ensure that self interested agents truthfully report their observations. Unfortunately, existing incentive compatible schemes have multiple Nash

equilibria. Moreover, lying equilibrium strategies usually yield higher payoffs than the truthful strategy.

In this paper we analyze the influence of trusted reports on the set of Nash equilibria of two existing incentive-compatible reputation mechanisms. We emphasize the existence of lying Nash equilibria, and investigate how such undesired equilibrium points can be eliminated. By having a fraction of trusted reports it is possible to have a mechanism where the truthful strategy is the only (or the most attractive) strategy to be followed.

Besides the numerical analysis of the two reputation mechanisms we also provide a general methodology for eliminating undesired equilibrium points from incentive compatible reputation mechanisms.

References

1. K. Aberer and Z. Despotovic. Managing Trust in a Peer-2-Peer Information System. In *Proceedings of the Ninth International Conference on Information and Knowledge Management (CIKM)*, 2001.
2. A. Birk. Learning to Trust. In R. Falcone, M. Singh, and Y.-H. Tan, editors, *Trust in Cyber-societies*, volume LNAI 2246, pages 133–144. Springer-Verlag, Berlin Heidelberg, 2001.
3. V. Conitzer and T. Sandholm. Complexity Results about Nash Equilibria. In *Proceedings of the IJCAI*, Acapulco, Mexico, 2003.
4. R. Cooke. *Experts in Uncertainty: Opinion and Subjective Probability in Science*. Oxford University Press: New York, 1991.
5. C. Dellarocas. Goodwill Hunting: An Economically Efficient Online Feedback. In J. Padget and et al., editors, *Agent-Mediated Electronic Commerce IV. Designing Mechanisms and Systems*, volume LNCS 2531, pages 238–252. Springer Verlag, 2002.
6. R. Jurca and B. Faltings. An Incentive-Compatible Reputation Mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.
7. R. Jurca and B. Faltings. “CONFESS”. An Incentive Compatible Reputation Mechanism for the Online Hotel Booking Industry. In *Proceedings of the IEEE Conference on E-Commerce*, San Diego, CA, USA, 2004.
8. D. M. Kreps, P. Milgrom, J. Roberts, and R. Wilson. Rational Cooperation in the Finitely Repeated Prisoner’s Dilemma. *Journal of Economic Theory*, 27:245–252, 1982.
9. N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Informative Feedback: The Peer-Prediction Method. Forthcoming in *Management Science*, 2005.
10. T. G. Papaioannou and G. D. Stamoulis. An Incentives’ Mechanism Promoting Truthful Feedback in Peer-to-Peer Systems. In *Proceedings of IEEE/ACM CC-GRID 2005*, 2005.
11. B. Yu and M. Singh. Detecting Deception in Reputation Management. In *Proceedings of the AAMAS*, Melbourne, Australia, 2003.