

Truthful Reputation Information in Electronic Markets without Independent Verification

EPFL Technical Report ID: IC/2004/08

Radu Jurca and Boi Faltings
Artificial Intelligence Laboratory (LIA),
Swiss Federal Institute of Technology (EPFL)
CH-1015 Ecublens, Switzerland
{radu.jurca, boi.faltings}@epfl.ch
<http://liawww.epfl.ch/>

January 15, 2004 (last modified: April 6, 2004)

Abstract

Reputation mechanisms offer an efficient way of building the necessary level of trust in electronic markets. In the absence of independent verification authorities that can reveal the true outcome of a transaction, market designers have to insure that it is in the best interest of the trading agents to report the behavior in transactions truthfully. As opposed to side-payment schemes that correlate a present report with future reports submitted about the same agent, we present a mechanism that discovers (in equilibrium) the true outcome of a transaction by analyzing the two reports coming from the agents involved in the exchange. For two long-run rational agents, we show that it is possible to design such a mechanism that makes cooperation a stable equilibrium.

1 Introduction

The availability of ubiquitous communication through the Internet is driving the migration of business transactions from direct contact between people to electronically mediated interactions. People interact electronically either through human-computer interfaces or even through programs representing humans, so-called agents. In either case, no physical interactions among entities occur and the systems are much more susceptible to fraud and deception.

Traditional methods to avoid cheating, involving strong cryptography and Trusted Third Parties (TTP's) that overlook every transaction, are very costly and sometimes even impossible to apply due to the complexity and heterogeneity of the environment. The maintenance of the TTP's

can incur substantial costs, and network communities often have a strong desire of being independent of any authorities, as illustrated by the successful P2P systems.

Reputation mechanisms offer a novel and efficient way of ensuring the necessary level of trust which is essential to the functioning of any market. They are based on the observation that agent strategies change when we consider that interactions are repeated: the other party will remember past cheating, and changes its terms of business accordingly in the future. In this case, the expected future gains due to future transactions in which the agent has a higher reputation can offset the loss incurred by not cheating in the present transaction. This effect can be amplified considerably if such reputation information is shared among a large population and thus multiplies the expected future gains made accessible by honest behavior.

Existing reputation mechanisms enjoy huge success. Systems such as eBay¹ or Amazon² implement successful reputation mechanisms which are partly credited for the businesses' success. Studies show that users seriously take into account the reputation of the seller when placing their bids in online auctions ([9]) and that despite the incentive of free ride, feedback is provided in more than half of the transactions on eBay ([18]).

The major challenge associated with designing reputation mechanisms is to ensure that truthful reports are gathered about the actual outcome of the transaction. In a typical e-commerce transaction, e.g. an exchange between a seller (he) and a buyer (she), the buyer is required to first pay and then wait for the purchased good to be shipped to the intended destination. While the payment of the buyer can be easily verified with the authority intermediating the transaction (e.g. the credit card company), it is very difficult to verify that the seller has indeed shipped the promised good. We start from the assumption that the outcome of the transaction (i.e. the seller has shipped or not the good) is only known to the parties involved. Any reputation mechanism will therefore have information that is distorted by the strategic interests of the reporters.

Most e-commerce environments do not make it rational for an agent to report the truth. The private information of a buyer for example, about the trustworthiness of a seller is often regarded as an asset which should not be freely shared. Paying for the buyer's reputation report could overcome this inconvenient, however, no guarantee can be offered that the information provided is also true.

Incentive compatibility can be assured if the side payment for a reputation report is conditioned on the correlation with future reputation reports (assumed to be true) submitted about the same seller. [16] and [10] describe such schemes that make truth revelation a Nash equilibrium. A problem with these schemes however, is that they require certain constraints on the behavior of the sellers and on the beliefs of the reporting buyers: i.e. sellers have typed behavior, and the set of seller types to which buyers assign positive probability is countable and contains at least 2 elements. Moreover, such schemes are vulnerable to the collusion of reporting agents, they also make lying be a Nash equilibrium, and they

¹www.ebay.com

²www.amazon.com

are not robust against irrational buyers who lie from time to time.

In this paper we address the problem of honest feedback elicitation in a more general setting in which sellers and buyers are assumed to behave rationally. We base our findings on the key assumption that buyers also have a persistent presence in the market. Even though most of the theoretical models proposed by the academic community studying reputation mechanisms assume single shot buyers, we believe that it is more natural to consider that buyers also keep returning to the same business partners (sellers) during their lifetime. Human buyers definitely have this characteristic and therefore software agents that act on behalf of humans should also be modeled in a context of repeated interaction.

The single-shot buyer behavior is mainly motivated by the technical difficulties associated with maintaining a persistent identity while trading over multiple markets. As we will later show, our mechanism motivates buyers to maintain their identity, thus further motivating the validity of our long-run buyer assumption in the present context. Moreover, projects like the "Liberty Alliance" [14] address the technical problems by attempting to build a unified cross-market online identity which will make it even easier for the buyers to be recognized in different markets.

The idea behind our mechanism is that long-run buyers can obtain a better business deal if they develop a reputation for honestly reporting the outcome of the transactions. Rational sellers will fear cheating on buyers that have an established reputation as honest reporters (because the resulting negative report will be believed by the community and will affect the future revenues of the seller) and therefore a cooperative equilibrium can be achieved in which the reputable buyer obtains a better payoff than the rational buyer who submits the reputation reports that maximize a momentary side payment function.

More concretely, we propose a side payment scheme and a decision rule which determines the outcome of a transaction by correlating the two binary reports (*cooperation* or *defection*) coming from the buyer and the seller involved in that transaction. For one round, three cases are possible:

1. The seller admits having defected. Regardless of the buyer's report, we can conclude in this case that the seller indeed defected. For a seller, falsely acknowledging defection implies a double loss (i.e. the future loss due to a negative reputation report, and the momentary loss coming from not taking the opportunity of defecting) and therefore no rational seller will report *defection* without actually defecting.
2. Both agents report *cooperation*. In this case a cooperative outcome can be recorded for the transaction in question if we make it impossible for the seller to bribe the buyer into untruthfully submitting a cooperative report.
3. The seller claims *cooperation* while the buyer reports *defection*. In this case, we can only be sure that one of the agents is lying. Since untruthful reporting is what we seek to avoid, both agents will be punished in this case: a negative report is being recorded for the seller, and both the buyer and the seller are fined for lying.

By following the above described protocol, we show that for a long-run buyer it is more profitable in some circumstances to report the truth, even when that means withstanding the momentary fine due to a situation in which the seller cheated but reported cooperation. This apparently irrational behavior of the buyer determines the seller to change his future behavior due to the threat that every unconfessed defection will be met with hostility by a buyer who probably always reports the truth. In equilibrium, we show that there is a finite upper bound on the number of times a rational seller is willing to defect and also report cooperation, which in turn, determines an upper bound on the number of non-cooperative transactions between two long-run agents.

Section 2 presents the related work, Section 3 describes the assumptions that we make about the environment, the mechanism itself, and a game-theoretic analysis of the repeated interaction between a seller and a buyer using our mechanism. Section 4 presents some open issues of the presented mechanism and directions for future work. Finally, Section 5 concludes our work.

2 Related Work

The notion of trust is used to refer to a subjective decision making process that takes into consideration a diversity of factors. The *Social Auditor Model* [11] is one of the existing models that explain how humans take trust decisions by using a set of rules. One of the input information that is often used in a trust decision making process is the *reputation* of the partner. Reputation can be regarded as a unitary appreciation of the personal attributes of the trustor: competence, benevolence, integrity and predictability. [17] presents an extensive classification of reputation by the means of collecting it.

Theoretic research on Reputation Mechanisms started with the three seminal papers of Kreps, Milgrom, Wilson and Roberts [12, 13, 15] who introduced the *reputation effect*, i.e. preference of agents to develop a reputation for a certain "type". A *type* is an apparently ³ irrational behavior that obeys some exact rules: e.g. an agent that cooperates all the time in a repeated Prisoners' Dilemma game is said to have a cooperative type. If a player (player one) is convinced that her opponent (player two) has a certain type, she will divert from playing the equilibrium strategy to playing a best-response strategy against the opponent's type. This new equilibrium might give player two a higher payoff than the initial one; it is therefore rational for player two to build a reputation for a certain type (commitment type) in order to eventually convince player one to play a best response strategy against her commitment type. Building a reputation involves some costs (as player one might not easily accept to revert to a best-response strategy) which have to be outrun by the future payoffs obtained when the reputation becomes credible. As a consequence,

³Typed behavior can be rationally explained by the existence of different payoff matrixes for the players of the same game. What seems rational for one agent (with one payoff matrix) might seem totally irrational for another agent having a different payoff matrix

the reputation effect exists only in a certain class of games, with players meeting certain criteria.

Fudenberg and Levine [7] study the class of all repeated games in which a long-run player faces a sequence of single-shot opponents who can observe all previous games. Based on the reputation effect, if the long-run player is sufficiently patient and the single-shot players have a positive prior belief that the long-run player might be a commitment type, the authors derive a lower bound on the payoff received by the long-run player in any Nash equilibrium of the repeated game. This result holds for both finitely and infinitely repeated games, and it is robust against further perturbations of the information structure (i.e. it is independent of what other types may exist with positive probability).

Schmidt [20] provides a generalization of the above result for the two long-run player case in a special class of games called of "conflicting interests", when one of the players is sufficiently more patient than the opponent. A game is of conflicting interests when the commitment strategy of one player (player one) holds the opponent (player two) to his minimax payoff – [20], Definition (1). The author derives an upper limit on the number of rounds player two will not play a best response to player one's commitment type, which in turn generates a lower bound on player one's equilibrium payoff.

[1, 2] describe computational trust mechanisms based on direct interaction-derived reputation. Agents learn to trust their partners, which increases the global efficiency of the market. However, the time needed to build the reputation information prohibits the use of this kind of mechanisms in a large scale online market.

A number of reputation mechanisms also take into consideration indirect reputation information, i.e. information reported by peers. [19, 21, 22] use social networks in order to obtain the reputation of an unknown agent. Agents ask their friends, who in turn can ask their friends about the trustworthiness of an unknown agent. Recommendations are afterwards aggregated into a single measure of the agent's reputation. This class of mechanisms, however intuitive, does not provide any rational participation incentives for the agents. Moreover, there is little protection against untruthful reporting, and no guarantee that the mechanism cannot be manipulated by a malicious provider in order to obtain higher payoffs.

Dellarocas [5] presents an efficient binary reputation mechanism that encourages a cooperative equilibrium in an environment of purely rational buyers and sellers. The mechanism is centralized, it works for single-value transactions, however, the buyers do not have any incentives to provide feedback.

A decentralized reputation mechanism is presented in [10]. Reputation feedback is collected, aggregated and disseminated by some specialized agents who are independent. For agents whose behavior can be modeled by a "dynamic type", the authors describe a side payment scheme that makes it rational for agents to truthfully report their observations.

In the same group of work that addresses the necessary property of incentive compatibility, we mention [3, 4, 16]. [3] considers exchanges of goods for money and proves that a market in which agents are trusted to the degree they deserve to be trusted is equally efficient as a mar-

ket with complete trustworthiness. By scaling the amount of the traded product, the authors prove that it is possible to make it rational for sellers to truthfully declare their trustworthiness. Truthful declaration of one's trustworthiness eliminates the need of reputation mechanisms and significantly reduces the cost of trust management. However, the assumptions made about the trading environment (i.e. the form of the cost function and the selling price which is supposed to be smaller than the marginal cost) are not common in most electronic markets.

For e-Bay like auctions, the Goodwill Hunting mechanism [4] provides a way in which the sellers can be made 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 who has the power to modify the announcement of the seller. To our knowledge this is the best reputation mechanism for multi-value transactions.

A significant contribution towards eliciting honest reporting behavior is made in [16]. The authors propose scoring rules as payment functions which induce rational honest reporting. The scoring rules however, cannot be implemented without accurately knowing the parameters of the agents' behavior model, which can be a problem in real-world systems. Moreover, this mechanism can be used only when agents have typed behavior.

3 Reporting Truthful Reputation Information

3.1 Assumptions

We consider an environment in which the following assumptions hold:

- A rational seller interacts repeatedly with several rational buyers by trading one good of value v_i in each round i . The values $v_i \in (\underline{v}, \bar{v})$ are randomly distributed according to the probability distribution function ϕ ⁴;
- All transactions have a fixed profit margin equal to $(\rho_B + \rho_S)v_i$, where $\rho_S v_i$ is the profit of the seller and $\rho_B v_i$ is the profit of the corresponding buyer;
- All buyers are completely trustworthy: i.e. Each buyer first pays the seller and then waits for the seller to ship the good. The seller may defect by not shipping the promised good, and the buyer perfectly perceives the action of the seller;
- There is no independent verification authority in the market, i.e. the behavior of the seller in round i is known only to the seller himself and the buyer with which he traded in that round;
- The seller cannot refuse the interaction with a specific buyer, and can trade with several buyers in parallel. A buyer can however end

⁴Following the same argumentation proposed in [4], this model is valid for settings where the act of accumulating inventory is independent from that of (re)selling it: e.g. a highly dynamic used car dealership.

the interaction with the seller and choose to buy the goods from a completely trusted seller (e.g. a brick and mortar shop) for an extra cost representing a percentage (θ) of the value of the item bought. Once a buyer decides to terminate a business relationship with the seller, she will never trade again in this market. The seller, however, can always find other buyers to trade with.

- The buyer and the seller discount future revenues by δ_B and δ_S respectively. The discount factors also reflect the probability with which the agents are going to participate to the next transaction. $0 < \delta_S, \delta_B < 1$, and $\delta_S \gg \delta_B$ modeling the fact that the seller is likely to have a longer presence in the market than the buyer.
- The buyer and seller interact in a market (possibly a different one for each transaction) capable of charging listing fees and participation taxes.
- At the end of every transaction, both the seller and the buyer are asked to submit a binary report about the seller's behavior: a positive report, $R+$, signals cooperation while a negative report, $R-$, signals defection;

We also assume that in our environment there is a *semantically well defined, efficient* Reputation Mechanism (RM). Reputation is semantically well defined when buyers have exact rules for aggregating feedback into reputation information and for making trust decisions based on that reputation information. These rules determine sellers to assign a value to a reputation report ($R+$ or $R-$), reflecting the influence of that report on future revenues. RM is efficient if the values associated by sellers to reputation reports are such that in any transaction the seller prefers to cooperate rather than defect. If $V(R+, v)$ and $V(R-, v)$ are the values associated by the seller to the positive respectively the negative reputation report generated after a transaction of value v , we have: $V(R+, v) + \text{Payoff}(\text{cooperate}, v) > V(R-, v) + \text{Payoff}(\text{defect}, v)$ ⁵. A simple escrow service or Dellarocas' Goodwill Hunting Mechanism [4] satisfy these properties.

As the influence of reputation on the seller's future revenues can be isolated into a concrete value for each reputation report, every interaction between a seller and a particular buyer can be strategically isolated and considered independently. A rational seller will maximize his revenues in each such isolated interaction.

When perfect feedback (i.e. true and accurate) is available, a *well-defined, efficient* RM is enough to make rational sellers cooperate. Unfortunately, perfect feedback cannot be assumed. In the absence of independent verification means, we can only rely on the subjective reports submitted by the agents involved in the transaction; reports which are obviously biased by the strategic interests of the agents.

In the rest of this section we describe a mechanism that in equilibrium obtains true feedback about the outcome of the transaction by correlating the seller's and buyer's reports about that transaction.

⁵as an abuse of notation, we will sometimes use $V(R+, v) = V(R+)$ and ignore the fact that the value of a reputation report also depends on the value of the good.

3.2 The Mechanism

Every round i , a seller offers for sale a good of value v_i . The market charges the seller a listing fee ε_S , and advertises the good to the buyer. The buyer pays a participation tax ε_B , to the market, and the price v_i to the seller. If the seller cooperates, he ships the good directly to the buyer; otherwise the seller keeps the payment for himself and does not ship the good. After a certain deadline, the transaction is considered as over, and the market starts collecting information about the behavior of the seller. The seller is first required to submit a report. If the seller admits having defected, a negative report ($R-$) is submitted to the RM, the listing fees ε_S and ε_B are returned to the rightful owners, and the protocol is terminated. If, however, the seller claims to have cooperated, the buyer is also asked to provide a report. At this moment, the buyer can report cooperation, report defection, or she can report defection and terminate the interaction with the seller.

If the buyer reports cooperation, a positive reputation report ($R+$) is submitted to the RM, and the listing fees ε_S and ε_B are returned. If the buyer reports defection, both players will be punished as one of them is surely lying: a negative report ($R-$) is submitted to RM, and the listing fees ε_S and ε_B are confiscated. Finally, if the buyer decides to terminate the interaction, a negative report ($R-$) is submitted to RM, and the fees ε_S and ε_B are confiscated. Figure 1 provides a schematic description of the trading protocol of each round, i .

From game theoretic point of view, the above described protocol can be modeled by the extensive-form game $G = (N, (A_i), (\succsim_i), T)$, shown in Figure 2. $N = \{S, B\}$ is the set of players, the seller and the buyer respectively, $A_S = \{Cc_S, Cd_S, Dc_S, Dd_S\}$ is the action set of the seller, $A_B = \{c_B, d_B\}$ is the action set of the buyer, \succsim_S is the preference relation of the seller over the set of possible outcomes, \succsim_B is the preference relation of the buyer over the set of possible outcomes, and T is the player function, or the “turn” function which prescribes which player should make the next move after every possible game history. Let also \mathcal{A}_S and \mathcal{A}_B denote the set of all mixed strategies in G . An action profile a is a tuple (a_S, a_B) such that $a_S \in A_S$ and $a_B \in A_B$.

The outcome for the buyer is indicated as a single real value representing the buyer’s payoff in the current round. The outcome for the seller is indicated as a tuple $(X; P)$, where $X \in \{R+, R-\}$ represents the filed reputation report (positive or negative), and $P \in \mathbb{R}$ is the payoff obtained by the seller in the current transaction. The buyer’s preference relation \succsim_B is the “ \geq ” relation over the set of real numbers. We assume that the seller’s preference relation \succsim_S has the following properties:

1. $(X; P_1) \succsim_S (X; P_2)$ for $X \in \{R+, R-\}$ and $P_1 \geq P_2$
2. $(R+; \rho_S v_i) \succsim_S (R-; (1 + \rho_S) v_i)$ for any v_i

While the first property is driven by common sense, the second property is guaranteed by the efficient reputation mechanism present in the market.

We further assume that the preference relations of the players can be described by a payoff function. Let $o_i(a_S, a_B)$ be the outcome for

1. The seller offers for sale a good of value v_i .
2. The market charges the seller a listing fee ε_S and posts the product for sale. ε_S is the lying fine imposed by the market to the seller if contradictory reports are submitted.
3. The buyer pays v_i to the seller and the tax ε_B to the market. ε_B is the lying fine imposed by the market to the buyer if contradictory reports are submitted.
4. The seller decides whether or not to ship the good (i.e. whether to cooperate or defect). If the seller cooperates, he ships the good directly to the buyer.
5. The market requests the seller to submit a binary report (c_S for cooperation or d_S for defection) about his own behavior in the current round.
6. If the seller reports d_S , a negative report $R-$ is sent to the RM, and the market pays ε_B to the buyer and ε_S to the seller. The transaction is completed.
7. If the seller reports c_S , the market asks the buyer to submit a report. The buyer can report cooperation (c_B), defection (d_B) or she can quit the game (*out*).
8. If the buyer reports c_B , a positive report $R+$ is sent to the RM, and the market pays ε_S to the seller and ε_B to the buyer. The transaction is completed.
9. If the buyer reports d_B , a negative report $R-$ is sent to the RM, and the market pays nothing to either the seller or the buyer.
10. If the buyer decides to quit the game, a negative report $R-$ is sent to the RM, and the market pays nothing to either the seller or the buyer.

Figure 1: Description of the transaction protocol.

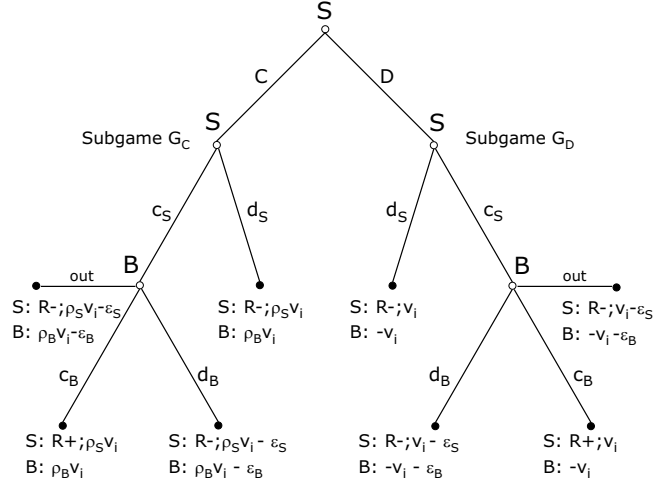


Figure 2: Game G modeling the one-round interaction protocol.

player $i \in \{S, B\}$ in G , when the seller plays $a_S \in A_S$ and the buyer plays $a_B \in A_B$. The payoff function g_i of player i maps outcomes to real number payoffs: $g_i[o_i(a_S, a_B)] \in \mathbb{R}$ represents the payoff for player i corresponding to the outcome o_i . A payoff profile $v(a)$ corresponding to the action profile $a = (a_S, a_B)$ is the tuple (v_S, v_B) such that $v_S = g_S(o_S(a_S, a_B))$ and $v_B = g_B(o_B(a_S, a_B))$.

For the buyer, the mapping from outcomes to payoffs is straightforward. For the seller, we assume that:

- $g_S[(X; P)] = V(X) + P$ for any $X \in \{R+, R-\}$ and $P \in \mathbb{R}$;
- $g_S[(R+; \rho_S v_i)] = g_S[(R-; (1 + \rho_S)v_i)] + \epsilon$ where $\epsilon > 0$, according to the second property of the seller's preference relation \succ_S .

The repeated transaction between the seller and one buyer can be modeled by a T -fold repetition of the stage game G , denoted G^T , where T might be finite or infinite. In this paper we will deal with the infinite horizon case, however, the results obtained can be applicable with minor modifications to finitely repeated games as well if T is large enough.

In the repeated game, player i obtains the average discounted payoff:

$$V_i = (1 - \delta_i) \sum_{\tau=0}^T \delta_i^\tau g_i^\tau; \quad (1)$$

where δ_i denotes her (his) discount factor, and g_i^τ is the payoff obtained by player i in round τ .

We define the *average continuation payoff* for player i from period t onward (and including period t) as:

$$V_i^t = (1 - \delta_i) \sum_{\tau=t}^T \delta_i^{\tau-t} g_i^\tau; \quad (2)$$

After each round, both players perfectly perceive the action of the opponent. They have perfect recall and can condition their play on the entire past history of the game. We denote by h^t a specific history of the repeated game out of the set $H^t = (A_S \times A_B)^t$ of all possible histories up to and including period t . A pure strategy s_i of player i in the repeated game is a sequence of maps $s_i^t : H^{t-1} \rightarrow A_i$. Correspondingly, let σ_i denote a mixed strategy of player i , where $\sigma_i^t : H^{t-1} \rightarrow \mathcal{A}_i$. By $V_i^t(\sigma_S, \sigma_B)$ we denote the overall payoff for player i obtained from period t onward (and including period t) if the seller follows the strategy σ_S and the buyer follows the strategy σ_B .

3.3 Equilibrium Analysis

For discounted infinitely repeated games with perfect information, the Folk Theorem [8] guarantees that every enforceable outcome (i.e. feasible and individually rational) can be obtained by a subgame perfect equilibrium (SPE) strategy profile when the discount factors are big enough. The results of this theorem do not apply directly to the game G^∞ because in every round t we allow the buyer to quit the game.

When the buyer terminates an interaction with a seller (chooses *out* in round t), she obtains a continuation payoff equal to:

$$\hat{V}_B^{t+1} = (1 - \delta_B) \sum_{\tau=t+1}^{\infty} \delta_B^{\tau-t-1} v_\tau (\rho_B - \theta);$$

If we denote by \tilde{v} the average value of a transaction, the expected value of \hat{V}_B^{t+1} is:

$$E[\hat{V}_B^{t+1}] = \tilde{v}(\rho_B - \theta);$$

Any SPE strategy profile must give the buyer at least \hat{V}_B^{t+1} after every round t (otherwise the buyer can profitably deviate to *out* in round t). The minimum continuation payoff of the buyer is therefore:

$$\underline{V}_B^t = (1 - \delta_B)(-v_t - \varepsilon_B) + \delta_B \hat{V}_B^{t+1}; \quad (3)$$

A payoff profile $\hat{v} = (\hat{v}_S, \hat{v}_B)$ dominates another payoff profile $v = (v_S, v_B)$ if it is better for at least one of the players and not worse for any of the players: i.e. there is $i \in \{S, B\}$ such that $\hat{v}_i > v_i$ and for all $j \in \{S, B\} \setminus i$, $\hat{v}_j \geq v_j$.

We restrict our attention to SPE strategies of G^∞ which are not dominated. A SPE strategy s is not dominated if there is no other SPE strategy \hat{s} such that the the payoff profile generated by \hat{s} dominates the payoff profile generated by s in G^∞ . The intuition behind this assumption is that no player will choose to play a SPE strategy as long as there is another SPE strategy which can bring him a higher payoff while not decreasing the payoff of the opponent.

The above restriction limits the set of SPE strategies to the ones generating an equilibrium path containing a mixture of the action profiles (C_{c_S}, c_B) and (D_{c_S}, c_B) .

Lemma 1 *All not dominated SPE strategies prescribe only the action profiles (Cc_S, c_B) and (Dc_S, c_B) on the equilibrium path in G^∞ .*

PROOF. Observe that any payoff profile in G is dominated by one of the payoff profiles corresponding to the action profiles (Cc_S, c_B) or (Dc_S, c_B) , i.e. $v(Cc_S, c_B)$ and $v(Dc_S, c_B)$ respectively.

Let s be a not dominated SPE strategy in G^∞ which prescribes for round t an action profile a other than (Cc_S, c_B) or (Dc_S, c_B) with positive probability. From s we construct strategy s' by replacing the action profile a in round t by the action profile $a' \in \{(Cc_S, c_B), (Dc_S, c_B)\}$ such that $v(a')$ dominates $v(a)$ in G . This replacement is possible because of the above observation. Moreover, the payoff generated by s' dominates the payoff generated by s .

We can show that s' is a SPE strategy in G^∞ by proving that there is no profitable one stage deviation for any of the players. Suppose that there is a profitable one stage deviation for player i from s' . Because of the way s' is constructed, any one stage deviation from s' is equally or less profitable than the corresponding one stage deviation from s . Therefore, the assumed one stage deviation will also be profitable for i in s . Contradiction. Thus s is dominated. ■

Let s be a mixed strategy profile such that with probability p the players play (Dc_S, c_B) and with probability $(1-p)$ the players play (Cc_S, c_B) . The expected continuation payoff of the buyer is:

$$\begin{aligned} E[V_B^{t+1}] &= E \left[(1 - \delta_B) \sum_{\tau=t}^{\infty} \delta_B^{\tau-t} [p(-v_\tau) + (1-p)\rho_B v_\tau] \right]; \\ &= \tilde{v}(\rho_B - p - \rho_B p); \end{aligned} \quad (4)$$

When playing in round t , the buyer knows which of the action profiles (Cc_S, c_B) or (Dc_S, c_B) are prescribed by the strategy s , and therefore the continuation payoff of the buyer is:

$$\begin{aligned} V_B^t|_{(Cc_S, c_B)} &= (1 - \delta_B)\rho_B v_t + \delta_B V_B^{t+1}; \\ V_B^t|_{(Dc_S, c_B)} &= (1 - \delta_B)(-v_t) + \delta_B V_B^{t+1}; \end{aligned} \quad (5)$$

depending on what s prescribes for round t . Since both $V_B^t|_{(Cc_S, c_B)}$ and $V_B^t|_{(Dc_S, c_B)}$ have to be greater or equal to \underline{V}_B^t , the maximum value of p is:

$$p \leq \bar{p} = \frac{(1 - \delta_B)\varepsilon_B + \delta_B \tilde{v}\theta}{\delta_B \tilde{v}(1 + \rho_B)}; \quad (6)$$

The upper bound on p limits the maximum attainable payoff, \bar{V}_S of the seller in G^∞ :

$$\bar{V}_S^t = (1 - \delta_S) \sum_{\tau=t}^{\infty} \delta_S^{\tau-t} [\bar{p}v_\tau + (1 - \bar{p})\rho_S v_\tau + V(R+)];$$

which has an expected value: $E[\bar{V}_S^t] = V(R+) + \bar{p}\tilde{v}(1 - \rho_S) + \tilde{v}\rho_S$. By replacing (6) we obtain:

$$\bar{V}_S^t = V(R+) + \tilde{v}\rho_S + \tilde{v}(1 - \rho_S) \frac{(1 - \delta_B)\varepsilon_B + \delta_B \tilde{v}\theta}{\delta_B \tilde{v}(1 + \rho_B)};$$

For any $p \in [0, \bar{p}]$ the strategy s can be made a SPE of G^∞ by adding minimax threats.⁶ Let us observe that when $p = 0$ our mechanism enforces the cooperative outcome and is incentive compatible. More precisely, our mechanism admits a SPE equilibrium in which the reputation mechanism collects only accurate feedback. Unfortunately, this equilibrium is not unique, and we can only guarantee that the maximum percentage of false reports is \bar{p} .

Following the ideas from [12], [7] and [20] we can limit the set of SPE strategies to a more desirable subset (i.e. consisting of those strategies which generate mainly true reputation reports and outcomes as close as possible to the socially efficient one) if we introduce a small amount of uncertainty in the perfect information game G^∞ .

A buyer who could commit to the “honest reporting” strategy, $s_B^* = (\text{play } c_B \text{ after } Cc_S \text{ and } d_B \text{ after } Dc_S)$, would benefit from cooperative trade. The seller’s best response against s_B^* is to play action Cc_S repeatedly, which leads the game into the socially efficient outcome. Unfortunately, under perfect information, the buyer’s commitment for s_B^* is not credible: when actually asked to play c_B or d_B a rational seller prefers to play c_B .

However, if the seller has incomplete information in G^∞ (i.e. he believes that he might be facing a buyer who prefers to play the commitment strategy s_B^*) we show that it is possible for a rational buyer to build a reputation for playing as the commitment type. When the reputation becomes credible, the seller is convinced that the opponent buyer is playing as if she were committed to playing s_B^* and therefore switches to the best response strategy against s_B^* , i.e. the cooperative equilibrium. As an effect of reputation building, the set of equilibrium points is reduced to a set of points which are close to the socially efficient one, and which generate truthful reputation reports on the equilibrium path.

Formally, imperfect information can be modeled by a perturbation of the complete information repeated game G^∞ such that in period 0 (before the first round of the game is played) the “type” of the buyer is drawn by nature out of a countable set $\Omega = \{\omega_0, \omega_1, \dots\}$ according to the probability measure μ . The buyer’s payoff now additionally depends on her type. We say that in the perturbed game $G^\infty(\mu)$ the seller has incomplete information because he is not sure about the true type of the buyer.

Two types from Ω have particular importance:

- The “normal” type of the buyer, denoted by ω_0 , is the rational buyer who has the payoffs presented in Figure 2.
- The “commitment” type of the buyer, denoted by ω^* , always prefers to play the commitment strategy s_B^* .

In Theorem 1 we give an upper bound k_S on the number of times the seller is willing to play the action Dc_S in $G^\infty(\mu)$, given that he always observes the commitment strategy played by the buyer.

⁶Theorem 1 in [8] explains how the strategy can be built

The intuition behind this result is the following. The seller's best response to the commitment type buyer is to always cooperate and report cooperation, i.e. (Cc_S) , which gives the commitment type buyer her maximum attainable payoff in $G^\infty(\mu)$, corresponding to the socially efficient outcome. The seller however would be better off by playing against the normal type buyer. As we have seen above, against the normal type buyer, the seller can get more than the cooperative outcome by randomizing between the (Cc_S, c_B) and (Dc_S, c_B) action profiles.

A normal type buyer can be distinguished from a commitment type buyer only if the seller plays Dc_S . In this situation, the normal type buyer prefers to play c_B , while the commitment type buyer prefers to play d_B . The normal type buyer could however simulate the strategy of a commitment type buyer in order to obtain the payoff of the latter (i.e. the cooperative outcome).

Because the cooperative strategy involves a loss for the seller (i.e. the potential loss of not being able to get the higher payoff that could be obtained against the normal buyer) the seller should not become "easily" convinced that he is playing against a commitment type buyer. The question is therefore, how long should the seller try to determine the true type of the buyer. Because every outcome (Dc_S, c_B) (i.e. the seller tests the type of the buyer and the buyer plays the commitment strategy) generates a loss for the seller, and because the seller cannot wait infinitely for future payoffs (the seller's discount factor is less than 1) it follows that at some point, if the seller always observes the commitment strategy being played by the buyer, he must give up trying to test the true type of the buyer, and accept playing a best response against the commitment type buyer.

Before we proceed, we restate an important lemma of Fudenberg and Levine [7] about statistical inference. The lemma proves that if ω^* has positive probability and if the seller observes s_B^* being played in every round, then there is a fixed finite upper bound on the number of time the seller will believe s_B^* is unlikely to be played. The intuition behind this result is the following: if the seller believes that s_B^* will be played in the next round with probability less than $\bar{\pi}$, every time he observes s_B^* he is slightly surprised and therefore will update his beliefs accordingly. Because the commitment type of the buyer chooses s_B^* with probability 1, while the seller expects the buyer to choose s_B^* with probability smaller than $\bar{\pi}$, it follows from Bayes' Law that the seller's update of his belief of facing a commitment type buyer, is strictly greater than 0. However, this cannot happen arbitrarily often because the updated probability of the commitment type cannot become bigger than 1. This gives an upper bound on the number of periods in which the seller may expect s_B^* to be played with probability less than $\bar{\pi}$. [20]

Formally, any (possibly mixed) strategy profile (σ_S, σ_B) induces a probability distribution π over the set of histories $(A_S \times A_B)^\infty \times \Omega$. Given a history h^{t-1} , let $\pi^t(s_B^*)$ be the probability attached by the seller to the event that the commitment strategy s_B^* is being played in period t . Since h^{t-1} is a random variable, so is $\pi^t(s_B^*)$. Fix any $\bar{\pi}$, $0 \leq \bar{\pi} < 1$ and consider any history h induced by (σ_S, σ_B) . Along this history, let $n(\pi^t(s_B^*) \leq \bar{\pi})$ be the number of random variables $\pi^t(s_B^*)$ for which $\pi^t(s_B^*) \leq \bar{\pi}$. Again, since h is a random variable, so is n .

Lemma 2 Let $0 \leq \bar{\pi} < 1$. Suppose $\mu(\omega^*) = \omega^*$ and that (σ_S, σ_B) are such that $\text{Prob}(h \in H^* | \omega^*) = 1$, where H^* is the set of all histories in which the buyer always plays s^* . Then:

$$\text{Prob} \left[n(\pi^t(s_B^*) \leq \bar{\pi}) > \frac{\ln \mu^*}{\ln \bar{\pi}} \mid h \in H^* \right] = 0.$$

Furthermore, for any infinite history h such that the truncated histories h_t all have positive probability and such that s_B^* is always played, $\mu(\omega^* | h_t)$ is nondecreasing in t .

PROOF. See Fudenberg and Levine, [7], Lemma 1 ■

This lemma does not prove that the seller will become convinced that he is facing a commitment type buyer. It simply proves that after a finite number of rounds the seller becomes convinced that the buyer is playing as if she were a commitment type.

Theorem 1 If:

1. the seller has incomplete information in G^∞ ,
2. the seller assigns positive probability to the prior beliefs that the buyer is a “commitment” type and a “normal” type. i.e. $\mu(\omega_0) > 0$, $\mu_0^* = \mu(\omega^*) > 0$ and $\mu(\omega_0) + \mu(\omega^*) = 1$;

Then there is a finite upper bound k_s on the number of times the seller plays Dc_S in G^∞ .

PROOF. In proving this theorem, we will first show that a rational seller does not choose action Dc_S in any round in which he believes that the buyer will play d_B in G_D with a probability greater than a certain threshold. Having this threshold, we use Lemma 2 to derive an upper bound on the number of rounds in which the seller might play Dc_S .

Let μ_t^* be the probability of the belief the seller has before round t that he is facing a commitment type buyer. Let also π_t be the probability assigned by the seller to the event that the buyer is going to play the commitment strategy s_B^* in round t , such that $\pi_t \geq \mu_t^*$. Let $V_S^t(\mu_t^*)$ denote the expected continuation payoff of the seller prior to round t .

In round t the seller has to choose between playing Cc_S or Dc_S (actions Dd_S and Cd_S are strictly dominated by Cc_S). When playing Cc_S , the seller expects with certainty to obtain the outcome (Cc_S, c_B) , however, since both the rational and the commitment type will play c_B in this situation, he does not get any information about the type of the buyer. His expected payoff in this case is:

$$E[V_S^t(\mu_t^*) | Cc_S] = (1 - \delta_S)g_S(R+; \rho_S v_t) + \delta_S E[V_S^{t+1}(\mu_t^*)] \quad (7)$$

If the seller chooses to play Dc_S , he expects with probability π_t that the buyer will play d_B and with probability $(1 - \pi_t)$ that the buyer will reveal to be a normal type who plays c_B . Against the normal type, the seller can expect a continuation payoff of maximum \bar{V}_S^{t+1} . Therefore:

$$\begin{aligned} E[V_S^t(\mu_t^*) | Dc_S] &\leq \pi_t [(1 - \delta_S)g_S(R-; v_t - \varepsilon_S) + \delta_S E[V_S^{t+1}(\mu_{t+1}^*)]] \\ &\quad + (1 - \pi_t) \cdot [(1 - \delta_S)g_S(R+; v_t) + \delta_S \bar{V}_S^{t+1}]; \end{aligned} \quad (8)$$

From Lemma 2 we know that $\mu_{t+1}^* = \mu_t^*/\pi_t$, the sequence μ_t^* being non-decreasing. When $\mu_t^* = 0$ the seller is convinced that the buyer is a normal type, when $\mu_t^* = 1$ the seller is convinced that the buyer is a commitment type. Moreover, since the seller can always increase probability μ_t^* to any value, it must be that:

$$\begin{aligned} V_S^{t+1}(0) &\leq \bar{V}_S^{t+1} \\ V_S^{t+1}(1) &= (1 - \delta_S) \sum_{\tau=t+1}^{\infty} \delta_S^{\tau-t-1} g_S(R+; \rho_S v_\tau) \\ E[V_S^t(\mu_1)] &< E[V_S^t(\mu_2)] \text{ for any } \mu_1 > \mu_2; \end{aligned} \quad (9)$$

A rational seller chooses Dc_S in round t only if $E[V_S^t(\mu_t^*)|Cc_S] < E[V_S^t(\mu_t^*)|Dc_S]$. By replacing (7), (8) and (9) we obtain:

$$\pi_t < \frac{(1 - \delta_S)\bar{v}(1 - \rho_S) + \delta_S\Phi}{(1 - \delta_S)[\bar{v}(1 - \rho_S) + \epsilon + \epsilon_S] + \delta_S\Phi} = \bar{\pi}; \quad (10)$$

where:

$$\Phi = E[V_S^{t+1}(0)] - E[V_S^{t+1}(1)] = \tilde{v}(1 - \rho_S) \frac{(1 - \delta_B)\epsilon_B + \delta_B\tilde{v}\theta}{\delta_B\tilde{v}(1 + \rho_B)};$$

The seller will not choose action Dc_S in any round t in which $\pi_t > \bar{\pi}$. From Lemma 2 we know that there is a finite number of rounds in which π_t can be less than $\bar{\pi}$ and as a consequence there is a finite number of rounds in which the seller might play Dc_S . This bound is given by:

$$k_S = \left\lceil \frac{\ln(\mu_0^*)}{\ln(\bar{\pi})} \right\rceil \quad (11)$$

and depends on $\bar{v}, \tilde{v}, \epsilon, \epsilon_S, \epsilon_B, \delta_S, \delta_B, \rho_S, \rho_B, \theta$ and μ_0^* . \blacksquare

The existence of k_S further reduces the possible equilibrium payoffs a buyer can get in $G^\infty(\mu)$. When a normal type buyer is asked to play the action profile (Dc_S, c_B) according to some SPE equilibrium strategy s , the buyer can deviate to playing d_B and mimic the commitment type buyer (i.e. build a reputation for honestly reporting the behavior of the seller). In the worst case, a normal type buyer who mimics the commitment type will have to play k_S times the (Dc_S, d_B) action profile (until the seller becomes convinced that the buyer is playing as if she were a commitment type) followed by an infinite sequence of (Cc_S, c_B) (played when the seller is best responding to the commitment type buyer). In this case, the continuation payoff of the normal type buyer is:

$$\begin{aligned} \underline{V}_B^t &= (1 - \delta_B) \left[(-v_t - \epsilon_B) + \delta_B \sum_{\tau=t+1}^{t+k_S-1} \delta_B^{\tau-t-1} (-v_\tau - \epsilon_B) \right. \\ &\quad \left. + \delta_B^{k_S} \sum_{\tau=t+k_S}^{\infty} \delta_B^{\tau-t-k_S} \rho_B v_\tau \right]; \end{aligned}$$

Any equilibrium strategy in $G^\infty(\mu)$ must guarantee the normal type buyer at least \underline{V}_B^t . Let us reconsider the strategy s from the perfect

information game G^∞ according to which the players play (Dc_S, c_B) with probability p and (Cc_S, c_B) with probability $1-p$. By imposing that both $V_B^t|_{(Cc_S, c_B)}$ and $V_B^t|_{(Dc_S, c_B)}$ (Equation (5)) be greater or equal to \underline{V}_B^t , the maximum value of p is:

$$p \leq \bar{p}' = \frac{(1 - \delta_B)\varepsilon_B + (\delta_B - \delta_B^{k_S})(\tilde{v} + \varepsilon_B + \tilde{v}\rho_b)}{\delta_B\tilde{v}(1 + \rho_B)}; \quad (12)$$

However, the constraints on p presented in Equation (6) remain valid, and therefore $p \leq \min(\bar{p}, \bar{p}')$.

Particular importance has the case in which $k_S = 1$. \bar{p}' becomes:

$$\bar{p}' = \frac{(1 - \delta_B)\varepsilon_B}{\delta_B\tilde{v}(1 + \rho_B)}; \quad (13)$$

and as ε_B can be any positive value, \bar{p}' will in the limit approach 0. In this situation, the reputation mechanism will receive false reputation reports with vanishing probability.

The result of Theorem 1 has to be interpreted as a worst case scenario. In real markets, sellers that already have a small predisposition to cooperate will defect fewer times. Moreover, the mechanism is self enforcing, in the sense that the more buyers act as commitment types, the higher will be the prior beliefs of the sellers that buyers will report truthfully, and therefore the easier it will be for the buyers to act as truthful reporters.

The following properties are also straightforward to derive as a direct consequence of Theorem 1:

Property 1 *The mechanism is bounded socially efficient.*

SKETCH OF PROOF. Because of the lost exchange, outcome (Dc_S, c_B) generates a cumulated social loss of $(\rho_S + \rho_B)v_i$ every time it occurs. The perfect information equilibrium involves a possibly infinite number of rounds in which (Dc_S, c_B) is played. By limiting the number of times the seller is playing action D , we also limit to a finite number (i.e. k_S) the rounds in which the exchange does not occur. The social loss is therefore bounded above by $k_S(\rho_S + \rho_B)\bar{v}$ ■

Property 2 *The mechanism is weakly budget balanced*

SKETCH OF PROOF. The net payment to the mechanism is non-negative as every time there is a disagreement concerning the two reputation reports, the center gets $\varepsilon_B + \varepsilon_S$. By introducing supplementary service fees, the mechanism can be easily transformed into one that yields profit to the market. ■

4 Open Issues

Further benefits can be obtained if the buyers' reputation as honest reporters is shared within the market. A buyer that has once built a reputation for truthfully reporting the seller's behavior will benefit from cooperative trade during her entire lifetime, without having to convince each

seller separately. Therefore the upper bound on the loss a buyer has to withstand in order to convince a seller that she is a commitment type, becomes an upper bound on the total loss a buyer has to withstand during her entire lifetime in the market. How to efficiently share the reputation of buyers within the market remains an open issue.

Correlated with this idea is the observation that buyers that use our mechanism are motivated to keep their identity. In generalized markets in which agents are encouraged to play both roles (e.g. a peer-2-peer file sharing market in which the fact that an agent acts only as "seller" can be interpreted as a strong indication of "double identity" with the intention of cheating) our mechanism also solves the problem signaled in [6] related to the ease with which agents can change their online identity. The price to pay for the new identity is the loss due to building a reputation as a honest reporter when acting as a buyer.

The mechanism can be criticized for being centralized. The market acts as a central authority by collecting listing fees from the seller and the buyer, by asking the reputation reports at the end of each transaction, and by reasoning about the outcome of the transaction. However, as the mechanism does not require any information to be transmitted from one round to another (the seller stores the reputation of the buyer) we could have the same seller and buyer interact in multiple markets (decentralized system) without having to relay on one single centralized institution.

Our mechanism is not robust against further perturbations of the information structure (i.e. other buyer types that can exist with positive probability). The presence of "crazy" buyer types for example, (i.e. buyers who have a preference (or are indifferent) to reporting defection after the seller cooperated and rightfully reported cooperation) and a particular set of beliefs of the seller, could determine him to also play action Cd_s from time to time. Such an equilibrium can be sustained by the threat that any deviation from the equilibrium will trigger in the "crazy" buyer a deviation to always denounce defection. The assumptions and prior beliefs that can sustain such equilibria are quite un-natural, and therefore, highly unlikely to occur in real situations. See [20] Section 3 for a discussion on how to build such equilibria.

One direction of future research is to study the behavior of the above mechanism when there is two-sided incomplete information: i.e. the buyer is also uncertain about the type of the seller. A seller type of particular importance would be the "greedy" seller type who always likes to keep the partner buyer to a continuation payoff arbitrarily close to 0. In this situation we expect to be able to find an upper bound k_B on the number of rounds in which a rational buyer would be willing to test the true type of the seller. The condition $k_S < k_B$ would impose the constraints on the parameters of the system for which the reputation effect will work in the favor of the buyer: i.e. the seller will give up first the "psychological" war and revert to a cooperative equilibrium.

A somehow related problem is the robustness to mistakes, or imperfect monitoring of the opponent's actions. A seller's defection by mistake in a situation in which it was not rational for a seller to defect will be interpreted by the buyer as evidence of the seller's irrational behavior. A mechanism that can deal with two-sided incomplete information will be

able to also address this issue.

Last, but not the least, we plan to adapt this truthful reporting mechanism for reputation mechanism that affect the value of future transactions. For such mechanisms the repeated interaction between a buyer and seller is much more complicated to model. A negative report submitted by a buyer at time t might lead to more beneficial trade for that buyer in the future (since the negative reputation report will attract a decrease in the price of future sold goods). Making it rational for the buyer to submit the true report involves a detailed understanding of the underlying reputation mechanism, the solution being most likely application dependent.

5 Conclusions

In this paper we describe a truth elicitation mechanism for two long-run rational buyer and seller. The mechanism assumes the existence of a market able to disseminate data and collect fees from both parties and we rely on an efficient reputation mechanism to make it rational for the seller to give up the momentary gain obtained from cheating in favor of a positive reputation report. In the absence of any independent verification, we describe a transaction protocol that correlates the reports coming from the trading seller and buyer in order to determine the correct outcome of the transaction. In equilibrium we show that our mechanism collects false reputation reports with vanishing probability. As a consequence, true reputation information is supplied to the reputation mechanism which therefore enforces a cooperative equilibrium in the market.

References

- [1] 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.
- [2] A. Biswas, S. Sen, and S. Debnath. Limiting Deception in a Group of Social Agents. *Applied Artificial Intelligence*, 14:785–797, 2000.
- [3] S. Braynov and T. Sandholm. Incentive Compatible Mechanism for Trust Revelation. In *Proceedings of the AAMAS*, Bologna, Italy, 2002.
- [4] 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.
- [5] C. Dellarocas. Efficiency and Robustness of Binary Feedback Mechanisms in Trading Environments with Moral Hazard. MIT Sloan Working Paper #4297-03, 2003.
- [6] E. Friedman and P. Resnick. The Social Cost of Cheap Pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

- [7] D. Fudenberg and D. Levine. Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica*, 57:759–778, 1989.
- [8] D. Fudenberg and E. Maskin. The Folk Theorem in Repeated Games with Discounting or Incomplete Information. *Econometrica*, 54(3):533–554, 1989.
- [9] D. Houser and J. Wooders. Reputation in Internet Auctions: Theory and Evidence from eBay. University of Arizona Working Paper #00-01, 2001.
- [10] R. Jurca and B. Faltings. An Incentive-Compatible Reputation Mechanism. In *Proceedings of the IEEE Conference on E-Commerce*, Newport Beach, CA, USA, 2003.
- [11] R. Kramer. Trust Rules for Trust Dilemmas: How Decision Makers Think and Act in the Shadow of Doubt. In R. Falcone, M. Singh, and Y.-H. Tan, editors, *Trust in Cyber-societies*, volume LNAI 2246, pages 9–26. Springer-Verlag, Berlin Heidelberg, 2001.
- [12] 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.
- [13] D. M. Kreps and R. Wilson. Reputation and Imperfect Information. *Journal of Economic Theory*, 27:253–279, 1982.
- [14] Liberty Alliance Project. www.projectliberty.org.
- [15] P. Milgrom and J. Roberts. Predation, Reputation and Entry Deterrence. *J. Econ. Theory*, 27:280–312, 1982.
- [16] N. Miller, P. Resnick, and R. Zeckhauser. Eliciting Honest Feedback in Electronic Markets. Working Paper, 2003.
- [17] L. Mui, A. Halberstadt, and M. Mohtashemi. Notions of Reputation in Multi-Agents Systems:A Review. In *Proceedings of the AAMAS*, Bologna, Italy, 2002.
- [18] P. Resnick and R. Zeckhauser. Trust Among Strangers in Electronic Transactions: Empirical Analysis of eBay’s Reputation System. In M. Baye, editor, *The Economics of the Internet and E-Commerce*, volume 11 of Advances in Applied Microeconomics. Elsevier Science, Amsterdam, 2002.
- [19] M. Schillo, P. Funk, and M. Rovatsos. Using Trust for Detecting Deceitful Agents in Artificial Societies. *Applied Artificial Intelligence*, 14:825–848, 2000.
- [20] K. M. Schmidt. Reputation and Equilibrium Characterization in Repeated Games with Conflicting Interests. *Econometrica*, 61:325–351, 1993.
- [21] B. Yu and M. Singh. An Evidential Model of Distributed Reputation Management. In *Proceedings of the AAMAS*, Bologna, Italy, 2002.
- [22] B. Yu and M. Singh. Detecting Deception in Reputation Management. In *Proceedings of the AAMAS*, Melbourne, Australia, 2003.