A model of indirect reputation assessment for adaptive buying agents in electronic markets

Kevin Regan and Robin Cohen
School of Computer Science
University of Waterloo

Abstract. In this paper, we present a model for designing buying agents in electronic marketplaces that adapt by adjusting decisions about which sellers to select for business, based on reputation ratings provided by other buying agents in the marketplace (known as indirect reputation information). The focus of our research is a method for effectively representing and employing this indirect reputation information. In particular, we address the case of buying agents providing deceptive information to other buyers, by having each buyer model not only the reputation of all sellers in the marketplace but also the reputation of each buyer. We also systematically account for differing standards between buyers, in assessing the reputation of sellers. Overall, the model presented here builds on a strong foundation of how best to model seller reputation but allows for a suitably cautious integration of a social aspect to the reputation modelling, towards improved purchasing decisions for buyers.

Keywords: Electronic Markets, Adaptive Agents, Reputation, Deception
Topic Areas: Adaptive business agents, Negotiation (trust)

1 Introduction

Designing adaptive business agents in electronic marketplaces requires the development of algorithms for these agents to reason about the sellers in the marketplace, towards making effective purchasing decisions. In order to elicit good behaviour in a market, one approach has been to introduce flexible trust based systems. The use of reputation for promoting trust has been given considerable attention in the context of centralized eBay-like markets [2, 4], peer-to-peer networks [5], and multi-agent systems [1, 8, 6, 11, 12].

The centralized reputation systems used by Amazon.com and eBay aggregate buyer feedback to form a reputation for each seller. The advantage of this approach is that each buyer has access to a seller's global reputation, which can potentially take into account the experiences of every buyer that a seller has dealt with. However, since the marketplace is not directly involved in the delivery of goods, there is no simple way to verify the accuracy of buyer feedback. At the other end of the spectrum are decentralized multi-agent systems which assess reputation and trust from the perspective of a single agent [8]. A buyer will model the reputation of each seller it has direct interactions with. Since the buyer is directly involved, there is no doubt about the outcome of transactions, however the buyer is limited to information about sellers it has already done business with.

This paper attempts to take advantage of the strengths of the multi-agent approach while providing an agent the possibility of learning about sellers they have not yet interacted with. This paper also draws inspiration from mechanism design which according to Varian: [9] uses the tools of economics and game theory to design *rules of interaction* for economic transactions that will, in principle, yield some desired outcome.

We present a system for buyers to share their subjective reputations with the goal of avoiding dishonest sellers. We use the terms *direct* reputation and *indirect* reputation to distinguish between the seller reputation a buyer gains from direct interaction and the seller reputation a buyer receives from another buyer. Our approach will be to begin with the single-agent model given by Tran and Cohen [8] and proceed to extend this model by providing a system by which agents can ask other agents for assessments of a seller's reputation.

The system presented by this paper is embedded in an electronic market with agents in fixed roles as either buyers or sellers. A buyer is free to purchase a good from any seller; however a seller may deceive a buyer by not delivering the promised good or by over-advertising the quality of the good. Purchases are negotiated with a three step process similar to that of Contract Net [7]. A buyer issues a request for particular good, the sellers who wish to provide this good make a bid at a certain price, and the buyer chooses between these bids and selects a seller.

One of the challenges in designing an effective model for adaptive business agents to make use of indirect reputation information is how to address the differing standards for goods employed by the various buying agents in the marketplace. The value each buyer places on an item (say, a piece of clothing) may be a matter of individual preference and as a result the reputations each buyer establishes for a seller will be highly subjective. However, one can imagine groups of buyers who share the same tastes and frequent the same sellers who will share similar reputations for these sellers. This perspective can be used to develop algorithms for buying agents that are sensitive to possible subjective differences.

Another challenge that arises in electronic marketplaces is trust. Sellers might over-advertise the quality of goods or buyers may lie when exchanging information about sellers. Our model accounts for possible deception by modelling the reputation of agents and avoiding decisions based on deceptive information. We also try to design a model that can be relatively efficient, reducing the need for every buyer to communicate with every other buyer and that encourages cooperation amongst buyers.

In the next section, we outline our proposed approach. This is followed by sections that illustrate the approach through examples, and contrast it with other models in the field, demonstrating the overall contributions of our research. As will be seen, our model goes beyond others that allow for indirect reputation by explicitly adjusting for systematic differences in subjectivity between buyers and by carefully acquiring information to detect deceptiveness in these buyers.

2 Model

We use the model presented by Tran and Cohen [8] as a starting point because it provides a way for the agent to partition sellers into a set the agent has deemed reputable, a set the agent has deemed disreputable and a set of sellers about which the agent is still unsure. The model can function as originally intended when the agent knows about the other sellers and can be naturally extended to include indirect reputation when an agent is unsure. We begin with a closer look at the original single agent model.

2.1 Direct Reputation

The buyer primarily uses two criteria when selecting a seller. The buyer will use the reputation of a potential seller paired with an estimation of the value of the good that will be purchased.

Definition 1. *Given a set S of sellers. We denote the reputation of a seller $s \in S$ as seen by a buyer b as $r_s^b \in (-1, 1)$.*

For most of the properties discussed in this paper the use of superscript denotes who holds the property while the subscript denotes who it refers to.

Definition 2. *Given a set of goods G, a set of prices P, $f : G \times P \times S \to \mathbb{R}$ is the estimated value function used by a buyer to assess the value of a good given the price and seller. We generally denote the estimated value function for a buyer b as $f^b(\cdot)$.*

For a seller that a buyer has no previous experience with, the reputation is initially set to zero and the estimated value is simply a function of the price of the good. It should be noted that because the initial reputation is significantly higher than the lowest possible reputation value, we have the potential problem of existing sellers with bad reputations re-entering the market using false pseudonyms. Some possible methods for eliminating this problem will be explained in the discussion section of this paper.

We use a reputation threshold Θ and a disreputation threshold θ to partition the set of sellers. Sellers for whom $r^b > \Theta$ are deemed reputable (R). Sellers for whom $r^b < \theta$ are deemed disreputable (DR), while the rest of the sellers are put into the set (?)[1] which the seller is unsure of. We can formally express this as follows

$$\forall\, s \in S \quad s \in \begin{cases} S_R^b & \text{if } r_s^b > \Theta \\ S_{DR}^b & \text{if } r_s^b < \theta \\ S_?^b & \text{otherwise} \end{cases} \tag{1}$$

[1] Tran and Cohen use the term non-reputable to refer to this set

The reputation of a seller is adjusted based on resulting value of a transaction v^b and a buyer's satisfaction threshold ϑ^b. When $v^b \geq \vartheta^b$, the buyer is satisfied and the seller's reputation r_s^b is increased by $\mu(1-r_s^b)$. When $v^b < \vartheta^b$, the buyer is unsatisfied and the seller's reputation is decreased by $\nu(1-r_s^b)$. By setting $\nu > \mu$ we can ensure that reputation will be difficult to earn and easy to lose. We can also set μ and ν to be proportional to the amount of the transaction. Tran and Cohen [8] use

$$\mu = \frac{v^b - \vartheta^b}{\Delta v^b}, \quad \text{where } \Delta v^b = v_{max} - v_{min} \tag{2}$$

This method of adjustment provides two benefits. It allows the increase in reputation to be proportional to the value of the transaction. Thus, an expensive suit that was not delivered can impact a sellers reputation far more than a pair of socks a consumer is unsatisfied with. The other benefit of this approach is that there is evidence that making reputation difficult to build and easy to tear down will discourage sellers from changing the value of their goods and allowing their reputation to oscillate between periods of building a reputation and periods in which they milk the reputation by over-advertising the quality of goods [2].

The buyer choses the seller with the highest estimated value $f(\cdot)$ from among the reputable sellers. The potential sellers who have been deemed disreputable are never purchased from and the sellers a buyer is unsure of are occasionally used to buy goods from. The buyer selects a seller from the set $S_? \cap S^p$ with some small probability ρ in order to explore new sellers.

2.2 Indirect Reputation

We now move beyond the model presented by Tran and Cohen [8] to provide an approach using indirect seller reputation provided by other buyers.

Consider the situation after a buyer b has made a request for a good and received bids from a set S^p of potential sellers. In some situations it may be beneficial for the buyer to ask a set of other buyers about the potential sellers. We refer to other buyers in this role as *advisors*. For each advisor $a \in A \subseteq B$ our buyer will maintain a reputation r_a and partitions A_R, $A_?$, and A_{DR} in the same manner as seller information is maintained.

To reduce communications overhead the buyer will not seek the help of advisors when there is a set of reputable potential sellers to choose from. When there are no potential reputable sellers available, a buyer will ask the non-disreputable advisors (i.e. those in the set $A_R^b \cup A_?^b$) about a set S^a of sellers. The set S^a of sellers is composed of all the potential sellers the buyer is unsure about as well as a set $S_!^a$ of sellers which the buyer has interacted with in the past which is taken from $S_R^b \cup S_{DR}^b$. $S_!^a$ will allow our buyer to assess how each advisor's standards differ and adjust in order to correct for them.

In response the buyer will receive a set of reputations for each seller from each buyer that can be represented in using a matrix where $b_1 \ldots b_m \in A$, $S_?^A = s_1 \ldots s_n \in S_?$ and $S_!^A = s'_1 \ldots s'_n \in S_R \cup S_{DR}$.

The advisor responses are combined to form a new reputation r_s^A for each seller. This new reputation is used to construct a set of reputable potential sellers (as in equation 1) from which the buyer can make a more informed purchase decision. The way in which the advisor responses are combined must take into account the differing subjective standards used by each advisor to assess reputation as well as the possibility of the advisor being untruthful or inaccurate.

2.3 Advisor Subjectivity

To address the differing standards of an advisor, the buyer looks for any systematic difference in reputation and adjusts for it. As previously mentioned, our buyer asks each advisor about a set of sellers $S_!^a$ that it already knows about through direct experience. We use this set of sellers to assess the similarity of the advisor to the buyer as follows

Definition 3. *For each advisor* $a \in A_R^b \cup A_?^b$ *and seller* $s \in S_!^a$ *we may calculate the reputation error* $\epsilon_s^a = r_s^a - r_s^b$

If similar criteria were being used (of the relative value of price and quality for the buyer) and the advisor was being honest, then the error ϵ_s^a would approach zero. However, if there is a systematic difference in the way an advisor determines reputation, then ϵ_s^a may be large, but would be remain fairly constant over different sellers.

Definition 4. *We denote the mean and standard deviation of the reputation error over a set of sellers as* $\bar{\epsilon}^a$ *and* σ^a *respectively.*

To quantify this notion of how large ϵ_s^a is and how it varies, we find the mean $\bar{\epsilon}^a$ and standard deviation σ^a of ϵ^a across sellers. If σ^a is small, there is a systematic difference in the reputations that a has given b and we can adjust for this difference as follows:

$$\forall\, s \in S^a, \;\; r_s^a \leftarrow r_s^a - \bar{\epsilon}^a \tag{3}$$

2.4 Advisor Deception

Our buyer will use the reputation held for each advisor to mitigate the effects of deceptive or inaccurate reputations given by an advisor. To avoid confusion between these two notions of reputation, we will occasionally refer to the reputation an advisor has about a seller as a prediction, since when this is information is passed on to the buyer and used as indirect reputation the advisors are, in a sense, making a prediction about the outcome of the buyer's purchase.

The responses from each of the advisors are combined so that the effect of dishonest sellers is minimized. However, each advisor is assumed to be honest until we find sufficient evidence of deception. It should be noted that we do not adopt the approach of weighing an advisor's predictions by the advisor's reputation ($r_a^b \cdot r_s^a$) that has been used by the Sporas system [12] and others

[5, 10]. The argument for our approach is that a until an advisor is no longer reputable, it is beneficial to fully consider their prediction (and not dilute it by some fractional weight).

We lessen the impact of dishonest sellers by maintaining reputations for each advisor and only use the predictions of the reputable advisors. We begin by finding the average over all the reputable advisors for each reputable seller.

Definition 5. *Given a seller s and a set of reputable advisors $A_R^b \subseteq A$, we denote the average prediction about s over all $a \in A_R^b$ as $\bar{r}_s^A$.*

An advisor with a high reputation who decides to lie about a particular seller can still have a large impact. This is particularly relevant since we assume all advisors are reputable until proven otherwise. To lessen the impact of reputable dishonest advisors we can choose to ignore predictions that are significantly different from that of the other reputable advisors. As a measure of significant difference we use the standard deviation of the prediction given by the reputable advisors, which we denote σ_s.

$$r_s^A \leftarrow \text{ avg } r_s^a \text{ over } a \in A_R^b \text{ where } |r_s^a - \bar{r}_s^A| < \sigma_s \tag{4}$$

It should be noted that after a purchase a buyer's reputation for *all* of the advisors is updated. An advisor's reputation can increase even if it was ignored when the seller was being chosen. In this way an advisor who fell below the reputable threshold can be redeemed. Following the purchase, the buyer will either be satisfied or unsatisfied with the true quality of the good based on our satisfaction threshold ϑ and the reputation of the seller will be adjusted. We also adjust the reputation of each as advisor, essentially, based on whether they were right or wrong. If a buyer predicts a seller to be reputable and a buyer is satisfied the advisor's reputation is increased. Likewise, if a buyer predicts a seller to be disreputable, but due to information from other advisors our buyer decides to purchase from this seller and is dissatisfied, the reputation for our advisor will increase. In the cases where the advisor's prediction does not match the buyer's satisfaction, the advisor's reputation will decrease. After the adjustment of advisor's reputation, the advisors can be partitioned into reputable, unsure and disreputable sets using criteria similar to that of equation 1. The buying agent can then avoid returning to advisors who have been deemed disreputable[2].

3 Example

In this example we have only two potential sellers (s_r and s_{dr}) among whom our buyer b must decide to buy a good. The seller s_r has never deceived a customer, while s_{dr} has lied to customers. However, our buyer b, has no experience with either seller and turns to a set of advisors (a_1, a_2, a_3, a_4) for more information. For the purposes of our example, a_1 (from $A_?^b$) turns out to be deceptive and

[2] If at some point all advisors are deemed disreputable our model will revert to the direct reputation model of Cohen and Tran [8]

provides deliberately inaccurate reputation information. a_2 is truthful, but has had good non-representative experiences with s_{dr} and provides an overly high reputation for this seller. Both a_2 and a_3 have high standards and this lowers the reputations they provide for each seller accordingly. a_4 is both truthful and has similar standards to our buyer.

Table 1. Example Details

	r_a^b	Explanation
a_1	-0.1	deceptive, and not trusted
a_2	0.4	truthful, high standards, inaccurate s_{dr} reputation
a_3	0.5	truthful, high standards
a_4	0.6	truthful, similar standards

	a_1	a_2	a_3	a_4
s_r	-0.2	-0.6	-0.7	0.2
s_{dr}	1.0	1.0	-1.0	-0.5

Now, our buyer b receives a reputation for s_r, s_{dr} and $s_! \in S_!$ from each advisor and if b were to simply average the reputations for s_r and s_{dr} without the methods developed to account for deception or differing standards, the result would be a reputation of -0.33 for s_r and 0.13 for s_{dr}. Now, let's say that b partitions sellers using: $\Theta = 0.20$ and $\theta = -0.20$ (as in equation 1), since $-0.33 < \theta$, s_r would be added to the set of disreputable sellers and since 0.13 is between θ and Θ, s_{dr} would be added to the set of sellers our buyer is unsure about.

The first step towards extracting accurate reputation information from our advisors is to account for any systematic bias. Our buyer finds the average difference between the reputation it holds and the reputation the advisor holds for each common seller $s'_i \in S_!$[3] In the case of a_2 and a_3, our buyer finds a difference of $\bar{\epsilon} = -1$ and a low σ indicating that our advisors consistently under-appreciate sellers by about -1. The buyer will adjust the reputations given by a_2 and a_3 by $-\bar{\epsilon}$. In our example 1 will be added to the reputations given by a_2 and a_3 and the average reputation for s_r and s_{dr} rises to -0.18 and 0.13 respectively[4].

The second step is to ignore any reputation information from advisors that our buyer is unsure about. Here, the buyer ignores the deceptively low reputation that a_1 provided for s_r and the deceptively high reputation that a_1 provided for s_{dr} resulting in s_r's reputation rising to 0.30 and s_{dr}'s reputation dropping to 0.17. The seller s_r is now in our buyer's reputable set, however our buyer is still unsure about s_{dr} due to the inaccurate high reputation given by the truthful advisor a_2.

The third and last step calculates the standard deviation of the set of reputations provided by reputable advisors and eliminates any reputation given by these reputable advisors that deviates from the average by more than one standard deviation. In our example the unrepresentative high reputation provided

[3] The reputation ratings for each s'_i held by the buyer and advisor are omitted here
[4] After adjustment a reputation greater than one will be normalized to one

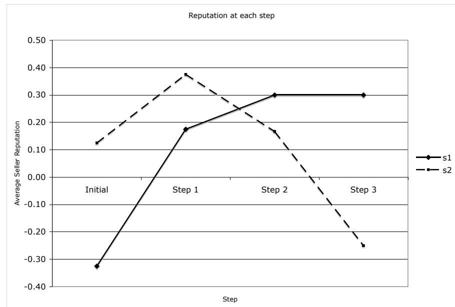

Fig. 1. average reputation at each example step

for s_{dr} by a_2 is eliminated and the resulting average reputation for s_{dr} drops to -0.25 moving s_{dr} into the set of disreputable sellers. In our example the methods developed in this paper have successfully limited the effect of differing standards, and deceptive and inaccurate advisors. The buyer selects s_r and the reputation of the advisors is adjusted, depending on whether the buyer is satisfied with the purchase. For example, the reputation of a_1 may be lowered, if the buyer was satisfied with s_r, and depending on the penalty applied, that advisor could be placed in A^b_{DR} and subsequently ignored.

4 Discussion

Our approach to modeling reputation in electronic marketplaces can be contrasted with that of other researchers. Yu and Singh present a similar approach to trust in multi-agent systems [10, 11]. An agent builds a reputation for correspondents with which it interacts. If an agent has had no previous contact with a correspondent, it seeks out other agents to act as witnesses relating the reputation have established about that correspondent. As with our model, the reputation given by the witnesses is adjusted based on how well the witness was able to predict the reputation of other correspondents. They do not, however, take into account the fact that witnesses may be using different standards in determining the reputation of sellers.

In their most recent work [11] an agent takes the set of transactions with a correspondent and assigns each transaction to one of two sets based on two thresholds (much in the way our Θ and θ are used to categorize sellers). A transaction above the first threshold is considered evidence for trustworthiness, while a transaction below the other threshold is considered evidence against. The reputation of a correspondent is essentially a 3-tuple with the number of transactions giving evidence for, against and neither. While this approach successfully captures the uncertainty in reputation and how uncertainty gives way with new evidence, the evidence for or against trustworthiness is not weighted by the value of a transaction.

Sabater and Sierra developed REGRET [6], a model of reputation that takes into account the personal, social and ontological aspects of reputation. This more complex reputation model goes beyond what we have addressed in this paper, but the way in which they address the personal and social aspects are similar to what we have done. In their model, the direct interactions between a buyer and a seller would build the buyers personal reputation of the seller. This personal reputation captures some phenomena that ours does not, such as the relevance of current interactions compared to those far in the past. The social aspect of reputation includes a method to combine reputations of a seller held by multiple buyers using weights to adjust for the differences between buyers. While REGRET provides a robust method for modeling many aspects of a seller, it does not address the possibility of dishonesty among agents sharing reputation information in the social aspect. In comparison our model offers a method to identify dishonest advisors in some cases even before making a purchase.

Zacharia, Moukas and Maes have proposed a collaborative reputation mechanism for electronic marketplaces called Sporos which assigns each user a reputation and allows for the ratings of a group of users to be combined to form the reputation. Like our approach, the reputation of other users is taken into account when forming this reputation; however, Sporos weights each rating by the reputation of the user. Sporos addresses the problem of cheap pseudonyms by initializing the reputation of each new user to the lowest possible reputation. While this discourages existing users from re-entering the market as new users, it also unduly penalizes new users. The paper also presents a novel model for using chains of trust in highly connected networks, but does not address *when* other agents should be consulted, *what* criteria should be used to find these other agents, or *how* they should address the subjectivity in each agent's ratings.

5 Conclusions and Future Work

The focus of this paper is to provide the potential for optimal market outcomes by reducing behavior which is detrimental to the welfare of the market as a whole. Specifically we examined how a system can be designed to limit the effect of deceptive sellers (and buyers) from the perspective of a single agent in a multi-agent system. The system developed allows the buyer to use indirect reputation gathered from other buyers acting as advisors to judge the reputation of sellers for which there is no direct reputation information. This model assumes that the indirect reputations provided by advisors is subjective and may not be truthful.

This paper offers two approaches to addressing deceptive advisors. The buyer will model the reputation of advisors and only listen to those who are deemed reputable. Also a buyer will ignore advisors whose predictions are significantly different from their peers in order to reduce the impact of deceptive advisors when combining indirect reputation gathered from a group of advisors. Since this process happens before a seller is selected, a deceptive advisor could be detected and dealt with before that advisor has had an opportunity to fool the buyer even once.

We have examined the challenges arising from subjectivity in reputations that are shared between buyers and we have offered an approach to identify any systematic bias in the reputations of sellers common to a pair of buyers and interpret future seller reputations to correct for this bias. By providing approaches to reduce the problems associated with subjectivity and the possibility of deception with advisors, this paper provides a solid foundation for use of indirect reputation to promote better market outcomes by reducing the impact of deceptive sellers. As a result, we lay the foundation for adaptive business agents to learn to avoid disreputable sellers by making use of reputation information provided by other buyers in the marketplace. We also provide a method for buying agents to assess the trustworthiness of both the sellers and their fellow buyers, towards making effective purchasing decisions.

There are several possible topics for future work, including possible theoretical extensions to the model and more detailed experimentation with the proposed approach. First of all, it would be useful to ensure that disreputable sellers do not re-enter the market using pseudonyms. Various strategies such as a trusted third party or the use of payments upon entry into the marketplace are discussed in [3]. The use of either of these approaches, while increasing the complexity of our system, would not conflict with any of the methods proposed for avoiding deceptive sellers and could be integrated into the overall framework, as future work. Another possible direction is to refine the model to find the best specific formulas for adjusting reputability, once an advisor's advice is evaluated. The formulas proposed in Tran and Cohen [8] are intended for updating a seller's reputation rating, but it may make sense to register disappointment in an advisor using a somewhat different adjustment factor.

It would also be worthwhile to further explore appropriate *rules of interaction* that encourage buying agents to be cooperative. These would work in conjunction with the incentives to be truthful that already exist in our model, with our buyers modeling the reputation of other buying agents and explicitly reasoning about possible deception. For example, an advisor cache could be added to store information about which advisors have helped a buyer in the past, to encourage the formation of small groups of buyers who are likely to reciprocate when asked for information about the reputation of sellers. In addition, there may be benefit in determining when agents are taking without giving, forming an asymmetric cycle of assistance amongst agents; isolating such cases may suggest revisions to groupings of agents formed to assist with purchasing decisions.

As for more experimental research, it would be useful to verify that the use of other buyers is beneficial when a buyer lacks reputation information about potential sellers. We can use the model provided by Tran and Cohen [8] as a base case since our approach reduces to their model when buyers do not communicate. To measure how far we have come towards meeting the challenges set out in section 3 the implementation should provide scenarios involving: randomly assigned deceptive sellers and buyers drawn from different distributions;

buyers with different levels of subjectivity regarding their reputations; deceptive sellers who attempt to shed bad reputations by re-entering the market as new users; increasing numbers of buyers, sellers and transactions to measure how the system scales. We can assess many of our design decisions by implementing them in isolation and observing their impact on the system as a whole. It will be useful to investigate how the implementation reacts to the adjustment of model parameters towards finding optimal values. It will be interesting to measure communication overhead to assess whether any benefits obtained through our approach outweigh the resulting increase in network load. We suspect in cases where buyers are new to the marketplace and lack experience about sellers that it will be possible to poll a series of other buying agents to obtain information that will indeed result in more effective purchasing decisions.

References

1. S. Breban and J. Vassileva. Using inter-agent trust relationships for efficient coalition formation. In *AI '02: Proceedings of the 15th Canadian Conference on Artificial Intelligence*, pages 221–236, 2002.
2. C. Dellarocas. Immunizing online repuation reporting systems against unfair rating and discriminatory behavior. *E-Commerce*, October 2000.
3. E. Friedman and P. Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10:173–199, 1999.
4. R. Jurca and B. Faltings. "confess": Eliciting honest feedback without independent verification authorities. *Sixth International Workshop on Agent Mediated Electronic Commerce (AMEC VI)*, 2004.
5. S. Marti and H. Garcia-Molina. Limited reputation sharing in p2p systems. In *EC '04: Proceedings of the 5th ACM conference on Electronic commerce*, pages 91–101, 2004.
6. J. Sabater and C. Sierra. Regret: reputation in gregarious societies. In *AGENTS '01: Proceedings of the fifth international conference on Autonomous agents*, pages 194–195, 2001.
7. R. G. Smith. The contract net protocol: high-level communication and control in a distributed problem solver. *Distributed Artificial Intelligence*, pages 357–366, 1988.
8. T. Tran and R. Cohen. Improving user satisfaction in agent-based electronic marketplaces by reputation modelling and adjustable product quality. In *AAMAS04*, pages 828–835, New York, USA, July 2004.
9. H. Varian. Economic mechanism design for computerized agents. In *Proceedings of the First USENIX Workshop on Electronic Commerce*, July 1995.
10. B. Yu and M. Singh. A social mechanism of reputation management in electronic communities. In *Proceedings of the 4th International Workshop on Cooperative Information Agents IV*, pages 154–165, 2000.
11. B. Yu and M. Singh. Detecting deception in reputation management. In *AAMAS03*, pages 73–80, 2003.
12. G. Zacharia, A. Moukas, and P. Maes. Collaborative reputation mechanisms in electronic marketplaces. In *32nd International Conference on System Sciences*, 1999.