Dynamic Pricing in a Reputation-Brokered Agent-Mediated Marketplace

Giorgos Zacharia, Theodoros Evgeniou¹, Pattie Maes

MIT Media Laboratory, Cambridge, MA 02139 USA

¹*MIT Artificial Intelligence Laboratory, Cambridge, MA 02139 USA
lysi@media.mit.edu, theos@ai.mit.edu, pattie@media.mit.edu*

We describe an Agent Mediated marketplace, with dynamically changing reputation ratings. In this marketplace, the seller reputations are updated in a collaborative fashion based on the performance of the user in the delegated tasks. We study the market with computer simulations of multiagent interactions, where sellers learn how to price their services dynamically. We first present some simple dynamic pricing methods and we investigate the different equilibria reached, based on the level of intelligence of the selling agents, the level of price-importance elasticity of the buying agents, and the level of unemployment in the marketplace. We then compare the equilibria reached with a theoretically "optimal" equilibrium that we show to exist. Based on the results of this comparison we design a new dynamic pricing algorithm that we experimentally show to be almost optimal for reputation-brokered agent mediated marketplaces.

Keywords: dynamic pricing, reputation mechanisms, multiagent systems

1 Introduction

This paper presents a framework for an Agent mediated Marketplace, where buyers are the users that need certain services that sellers can provide. In a marketplace of services we have to face complexities like measuring seller competency and performance - which can also often exist in marketplaces of tangible goods (i.e measuring the creditworthiness of the buyer, the delivery time of the seller, etc). We use collaborative reputation mechanisms to estimate the sellers' performance based on their past transactions, and we automate the processes of matchmaking and pricing of the services. The general framework of this marketplace has been introduced in [Zacharia, G. et al (2000)]. The buyers configure their agents with their budget and the importance of the specific task. These agents try to maximize their owners' utilities. In order to achieve this, the buyer agents estimate the expected performance of each seller based on the reputation of that seller in the marketplace and the seller's price, and choose the seller that maximizes their expected utility. Selling agents respond to buyers by bidding on behalf of their owners for the available tasks based on their owners' reputations. Due to the nature of our reputation mechanisms, there is an inherent market inefficiency [Friedman, E. and Resnick, P. (1998)], because initially sellers are undervalued until their reputation values come close to their actual ability. In order to solve, or at least minimize this problem, we suggest dynamic pricing algorithms. Dynamic pricing allows us to price as efficiently as possible by considering the current reputation of each seller.

In this paper we first show how, in the case of fixed sellers' reputations, simple dynamic pricing algorithms lead to equilibria that are not very different from the ones predicted by the theory. This is a finding similar to that of [Cliff, D. and Bruten, J. (1998)] in the sense that "zero intelligence plus" [Cliff, D. and Bruten, J. (1998)] is shown to lead to interesting market equilibria. We also show what market inefficiencies (compared with the theoretically optimal equilibrium) arise when these simple methods are used and the reputations change dynamically. We then present a new dynamic pricing algorithm that is experimentally shown to lead to an almost optimal equilibrium also when the reputations change dynamically.

2 Related Work

Kasbah [Chavez, A. and Maes, P. (1996); Moukas, A. et al (2000)] is an ongoing research project to help realize a fundamental transformation in the way people transact goods. In Kasbah, a user wanting to buy or sell a good creates an agent, gives it some strategic direction, and sends it off into the agent marketplace. Kasbah agents pro-actively seek out potential buyers or sellers and negotiate with them on their creator's behalf [Chavez, A. and Maes, P. (1996)]. In Kasbah, the reputation values of the individuals trying to buy/sell books/CDs are major parameters of the behavior of the buying, selling or finding agents of the system. The latest version of Kasbah, which is implemented using the MarketMaker [Wang, D. (1999)] infrastructure, allows users to trade intangible goods like translation services. The interface and the backend of our Knowledge Marketplace will also be implemented using MarketMaker but its major difference from Kasbah is the negotiation strategies of the buyers and sellers.

In Kasbah, the price negotiation is based on a limited number of predefined negotiation strategies provided by the system. Agents created with these strategies cannot adjust their negotiation behavior according to the market conditions and the user has to make sure that his/her price ranges are close to the market prices. In the Knowledge Marketplace, we wanted to have software agents automate the task of monitoring the market conditions for their users. This means that, instead of using predefined time-varying price functions we need to incorporate adaptive pricing for the sellers and utility evaluation functions for the buyers.

So far, there has been quite extensive research in adaptive pricing agents [Oliver, J. (1997); Cliff, D. and Bruten, J. (1998); Sairamesh J., and Kephart, J. O. (1998); Kephart, J., O., et al (2000)] and it was shown that with minimally intelligent agents (Zero Intelligent Plus) [Cliff, D. and Bruten, J. (1998)] we can achieve economically efficient equilibria without the agents knowing each other's strategy or the market conditions from a macroscopic level [Cliff, D. and Bruten, J. (1998)]. It was also shown that in a marketplace with quality differentiation and quality sensitive users, we can have stable price equilibria, as predicted by game theory [Sairamesh J., and Kephart, J. O. (1998)]. However, this was shown to be the case when the quality of the sellers is stationary and the sellers can sell their information goods to multiple buyers at the same time. In our experiments, we investigate the effects of dynamically changing reputations, for intangibles or services that engage the sellers to one buyer at a time.

3 Knowledge Marketplace

For the purposes of our model, we assume that the marketplace consists of the same number of buyers and sellers. The equilibria of this marketplace are evaluated for two different scenarios: unemployment (less demand than supply), and overemployment (more demand than supply). Since we keep the number of buyers and sellers fixed, we created the scenarios by changing the rate of creation of tasks for each buyer. In particular, the market operates in periods. In every period, each buyer has a probability P to generate a problem. Once a problem is generated, the buying agent dispatches a request for bids to all sellers. Upon receipt of this query, all available seller agents respond with a price bid and wait for the buyer's decision.

After the sellers respond to the buyer, the buyer evaluates the expected utility function for each bid and picks the available seller that offers the highest expected utility. Note that the buyer is allowed to reject all bids. Once it makes its selection, the buyer delegates the task to the chosen seller. The cycle repeats for all the buyers in the market. We repeat the simulation for a number of periods, and record all the contracts established, as well as the total utilities of the buyers and the total profits (sum of prices at which they managed to sell) of the sellers.

At every period each buyer can generate a problem of importance I with probability P . The importance I of the problem is a uniformly distributed random variable from 0 to 1. If a problem is generated, the buyer will request bids from the sellers without providing information about the importance of the task, so that it does not lose its bargaining power. The sellers, on the other hand, have uniformly distributed abilities A ranging from 0 to 1. The sellers reputation is updated over time based on the seller's ability. We discuss the reputation update mechanism below.

4 Reputation Mechanisms

Consumer to Consumer marketplaces like Kasbah [Chavez, A. and Maes, P. (1996)], MarketMaker [Wang, D. (1999)], eBay, Yahoo Auctions and Amazon Auctions introduce some major issues of trust [Kollock, P. (1999)]: Potential buyers have no physical access to the product or service of interest while they are bidding or negotiating. Therefore, sellers can easily misrepresent the condition or the quality of their products or services. Additionally, sellers or buyers may decide not to abide by the agreement reached at the electronic marketplace, asking later to renegotiate the price, or even refuse to commit the transaction. Even worse, they may receive the product or service and refuse to send the money for it, or the other way around. Almost all of these concerns are also true for marketplaces of intangible goods, except that instead of the uncertainty about the condition of the products we have the uncertainty about the competency or actual performance of the seller.

One way of solving the above mentioned problems would be to incorporate in the marketplace a reputation brokering mechanism, so that each user can customize his/her pricing strategies according to the risk implied by the reputation values of his/her potential counterparts. In our previous work [Zacharia, G. and Maes, P. (2000)], we have developed elaborate reputation mechanisms for open online marketplaces or communities that are robust against common abuses of online rating systems. A simple version of these reputation mechanisms is incorporated in our Knowledge Marketplace. For the full description of the reputation mechanism, we refer the reader to [Zacharia, G. and Maes, P. (2000)]. After a seller completes a task, his/her reputation will be updated, using the rating received from the buyer as an indication about his/her ability. Suppose that at time $t=i$, a user with reputation R_{i-1} is rated with a score W_i , which is a random value normally distributed around the user's ability A , truncated between 0 and 1. Let $E_i = R_{i-1}/D$, where D is the reputation range. At equilibrium, E_i can be interpreted as the expected value of W_i , which is the ability of the user A , though early in a user's activity it will be an underestimate. Let $\theta > 1$ be the effective number of ratings considered in our reputation evaluation [Zacharia, G. and Maes, P. (2000)]. We then propose the simplified Sporas formula [Equation 1], which is a recursive estimate of the reputation value of a user at time $t=i$, given the user's most recent reputation, R_{i-1} , and the rating W_i :

$$R_i = R_{i-1} + \frac{1}{\theta} \cdot \Phi(R_{i-1})(W_i - E_i) \quad , \quad \Phi(R_{i-1}) = 1 - \frac{1}{1 + e^{\frac{-(R_{i-1}-D)}{\sigma}}} \quad , \quad E_i = R_{i-1}/D$$

Equation 1 Sporas Formulae

The parameter σ controls the dumping function Φ so that the reputation of highly reputable users is less sensitive to rating fluctuations. In order for the agents to have no incentive to switch identities, we choose the initial reputation of the agents to be minimal. For the experiments described in this paper, we set the initial reputation to be 0.01.

5 Agent types

5.1 Buying Agents

In picking the most suitable seller for its task, the buyer tries to maximize the following Cobb-Douglas utility function:

$$U = (1 - P)^{1-l} O^l$$

Equation 2 Buyers' Utility Function

Where P is the price the buyer will pay, normalized by his budget cap ($P_{\text{actual}}/P_{\text{cap}}$) so that it is between 0 and 1, I is the importance of the problem to the buyer, and O is the outcome of the problem in the range $[0,1]$, where 1 is perfect and 0 is terrible. We chose this function for our simulation because it has the following properties:

1. For an important problem, the buyer is willing to spend more.
2. For an unimportant problem, the buyer will sacrifice quality for price.

Furthermore, we assume that a buyer always has the option to turn to some external market (for example, the traditional consulting market) with reputation 1, and price P_m to solve his problem. If none of the sellers' offers provides a greater utility to the buyer than the traditional market, then the buyer will employ the traditional market in solving his problem

In order to evaluate the expected utility, our buyers use the certainty equivalent principle. In other words, they treat the performance of the seller as a deterministic variable, represented by the value of his reputation. Thus they evaluate their utility function using Outcome = Reputation of the seller, which, as noted above, changes over time.

5.2 Selling Agents

There are several kinds of selling agents (sellers). First we present experiments with very simple pricing strategies: in this section we used Derivative Followers, Reputation Followers, and Random Sellers. In the following section we will discuss some more advanced dynamic pricing methods. Designing more intelligent sellers is the subject of ongoing research.

5.2.1 Derivative Followers

Derivative Followers (DF) are sellers who decide their next bid according to the success of the previous one. Therefore, these sellers focus on increasing their prices from one contract to another as long as they can get the contracts. Likewise, they decrease their bid when they do not get contracts. We assume that Derivative Followers increase their price bids by a fixed step S_{up} multiplied by a random number picked from a uniform distribution with range $[0,1]$ for the next ($inertia+I$) periods. The random number is different every time it offers a bid. Our preliminary experiments showed that the value of inertia does not have much effect on the results because there are no local maximums or minimums in the profit landscapes of the Derivative Follower sellers. Therefore, for this paper the inertia is set to 0. If a Derivative Follower fails to receive a contract, it will start decreasing its price bids by $S_{dn} * \text{random}$. In other words, if $idle$ is the number of periods after the inertia time passes, the Derivative Follower's offer will be:

$$P = \text{LastContractPrice} + S_{up} * \text{random}_1 - S_{dn} * \text{random}_2 * \text{idle}$$

Equation 3 Derivative Follower pricing

The random numbers are different for the two monomials, and they are both recomputed every time an offer is made. Such agents have also been used in other work (see for example [Sairamesh J., and Kephart, J. O. (1998)] and [Cliff, D. and Bruten, J. (1998)]).

5.2.2 Reputation Followers

Reputation Followers maintain a shadow price P_s on which they apply the Derivative Follower algorithm, assuming that they would offer that price if they had a perfect reputation. However, the price they actually announce is the product of the shadow price and the current reputation value, i.e. $P_s * R$. We decided to implement this scheme, in order to make agents who will actually respond fast to their reputation changes. Indeed, our experiments showed that Reputation Followers set bids that follow their reputation patterns and eventually their abilities better than the Derivative Followers do. In a sense our Reputation Followers are Derivative Followers but with a step that depends on their reputation and therefore is dynamically changing. Low reputation agents change their prices slowly; therefore, we expect that in the case of

unemployment they would initially perform better than low reputation Derivative Followers since they would undercut them. This was verified by our experiments.

5.2.3 Random Sellers

Random Sellers have no pricing or bidding strategies. They just bid random prices. Naturally, these agents are not expected to perform adequately, but they provide a measure to compare against. In section 7 we also compare the agents with ones that price according to a theoretically optimal strategy that we derive.

5.3 Sellers' Price Range Computation

Given reputation R , external market price P_m , and importance I , the maximum price that the seller can charge is the following:

$$P \leq 1 - \frac{(1 - P_m)}{R^{1/I}}$$

Equation 4 Maximum price for which a seller can expect to get a contract

Figure 1 The range of bids allowed for a seller, as his reputation increases, starting from 0.01, going to 1. Sellers have a chance of receiving the contract only if they bid below the curve given their reputation and the importance of the buyer's problem (which they do not know when placing the bid).

Furthermore, initially the sellers have a very low reputation ($=0.01$ in our case), therefore they can only receive low importance jobs. Even if they bid for zero price, they can only get a contract if:

$$R' \geq (1 - P_m)^{1-I} \leftrightarrow I \leq \frac{\log(1 - P_m)}{\log(R(1 - P_m))}$$

Equation 5 Price requirements relative to Reputation of Seller and Importance for the buyer making the evaluation of the seller's bid

where I is the importance of the job, and R is the initial reputation of the selling agent. This is expected, since agents will opt to build reputation, in order to be able to actively bid for a larger share of the contracts.

6 Results of Simulations

We ran several simulations to evaluate the behavior of our system and test the three simple pricing algorithms in the two market conditions. Without loss of generality, each seller also started with a minimal price, 0.1 (so that none of the sellers started with any advantage). The performance of the algorithms is evaluated based on the profits of each seller as a function of his/her ability. The three algorithms are also evaluated in competition settings where each third of the agents price with one of the algorithms. In this section, we present some preliminary results. The simulations were first run with many agents (100 or more) in order to get a general view of the agents' behavior. In section 7 we run experiments with a few agents to better track their behavior.

6.1 Unemployment

Figure 2 shows the profits of the sellers in the case that they are all either random, or Derivative Followers, or Reputation Followers (no competition among different pricing strategies). As shown, in the case of unemployment both followers perform better than Random Sellers, since Random Sellers often charge high prices even when they have low reputations, and therefore miss more contracts than the followers do. Regarding the two followers, we observe that when they do not compete both kinds of followers perform about the same on average: Figure 2 shows that for the particular simulation run Reputation Followers performed slightly better than Derivative Followers did (especially at the low ability range), but the difference is small.

On the other hand, when the three types of agents compete with each other, then all the agents with more than random intelligence were observed to drive their prices down in order to attract the agents of the buyers. Therefore Random Sellers were not able to get contracts so, as the experiments showed, almost all Random Sellers had zero profit. Furthermore, we observed that some followers also could not escape from their initial low reputation by offering low enough prices (that can be attributed to the randomization in following the derivative). Even some agents with very high abilities were not able to engage in trade and could not raise their reputations. Other agents that initially offered lower prices raised their initial reputation, and, thus, attracted even more buyers. This scenario is a good example of how initial history might affect such a marketplace with positive reputation mechanisms.

Furthermore, as shown in Figure 3, Reputation Followers tend to escape from their initial low reputation more often than other agents do (we get back to this issue in section 7 for the "no competition" case). This is because at the initial states they increase their prices slowly (since they have low reputation) therefore undercutting the Derivative Followers and the Random Sellers, and therefore accumulating more profit. In Figure 4 we show the average difference between the profits of the Reputation Followers and the Derivative Followers over time. As shown in the figure, at the beginning (when all agents have low reputations) the difference between the profits of the Reputation Followers and the Derivative Followers increases (Reputation Followers undercut the Derivative Followers most of the time). Over time this difference decreases (over time the reputation of the agents that managed to "escape" the minimum reputations is the same as their actual ability so both Reputation Followers and Derivative Followers behave similarly). The figure shows the difference of profits for two kinds of agents: low ability ones (in this case we chose the agents with ability less than 0.3 and measured their average profit) and high ability

ones (agents with ability larger than 0.7). The phenomenon appears for both types of agents, and it is stronger for the high ability ones.

Figure 2 Unemployment simulations: plus signs (+) are Random Sellers, stars signs (*) are Derivative Followers, and circle signs (o) are Reputation Followers. We run three experiments where for each one all sellers had the same strategy. The figure compares their final profits.

Figure 3 Unemployment simulations: plus signs (+) are Random Sellers, stars signs (*) are Derivative Followers, and circles signs (o) are Reputation Followers. We run one experiment where a third of the sellers where random, the other third Derivative Followers, and the remaining Reputation Followers.

Figure 4 Unemployment simulations: Average profit difference between reputation and Derivative Followers over time (y-axis: (avg. profit of rep. fol. – avg. profit of der. fol.) / (# of trade iterations), x-axis: # of trade iterations). Bold line for low ability sellers, thin line for high ability sellers.

6.2 Overemployment

We have overemployment when we expect that a seller will be "guaranteed" to get a job. This occurs when $p \cdot B > S$, where S is the number of sellers, B is the number of buyers, and p is the probability of a job being created by the buyer at a market period. If $p \cdot B > S$, then the seller can be employed continuously (without having a single period of unemployment), as long as his price offers fall within the acceptable range of the buyers.

In the case of overemployment, all the sellers have the potential of enormous profit. However, the Random Sellers do not behave that well (since they do now take advantage of the overemployment), as expected, while the Reputation Followers perform according to their abilities. In the case of overemployment, the Derivative Followers perform overall the best, as shown in Figure 5, Figure 6, and Figure 7 (which show the same results as Figure 2, Figure 3, and Figure 4 for the case of overemployment).

Figure 5 Overemployment simulations: convention as in Figure 2.

Figure 6 Overemployment simulations: convention as in Figure 4.

Figure 7 Overemployment simulations: convention as in Figure 5.

7 Comparison with theoretically optimal policies

In this section we present a derivation of a theoretically “social optimal” pricing strategy that leads to a pricing equilibrium. We then compare the agents above with the optimal sellers we derive. Finally we present a new almost optimal dynamic pricing method.

7.1 Optimal Pricing under Complete Information

Suppose we have n sellers and m buyers, and let us represent the set of our sellers as $\{S_1, S_2, S_3, \dots, S_n\}$, sorted by reputation such that:

$$R(S_1) > R(S_2) > R(S_3) > \dots > R(S_n)$$

and our buyers $\{B_1, B_2, B_3, \dots, B_m\}$, sorted by quality sensitivity such that

$$I(B_1) > I(B_2) > I(B_3) > \dots > I(B_m)$$

We have unemployment conditions when $n > m$, full employment when $n = m$, and overemployment when $n < m$.

Under any conditions the maximum number of trades that can take place in each round is: $t = \min(n, m)$.

We will now prove that *if there is perfect information in the marketplace, such that the sellers know each other's reputation, the buyers' utility functions, then there exists a Nash equilibrium [Nash, J. (1950)]. In the Nash equilibrium, the prices are such that the sellers and the buyers will pair according to their respective ability and quality sensitivity. Therefore, we will observe trades among the following pairs:*

$(S_1, B_1), (S_2, B_2) \dots (S_t, B_t)$

Proof:

We have a Nash equilibrium if for every consecutive pairs of sellers and buyers $(S_i, B_i), (S_j, B_j)$ such that $1 \leq i < j \leq t$ then:

The prices $P(S_i)$ and $P(S_j)$ offered by sellers S_i and S_j are such that if any of the sellers prices higher or lower then they will be worse off.

Given the prices $P(S_i)$ and $P(S_j)$, then the buyers B_i and B_j are both better off by selecting sellers S_i and S_j respectively.

Therefore seller S_i will price its services such that:

$$U_i(P(S_i), R(S_i), I(B_i)) > U_i(P(S_j), R(S_j), I(B_i))$$

Equation 6 Partial Utility ordering for successful seller bidding

otherwise B_i could have selected S_j (if B_i was choosing first), and

$$U_j(P(S_i), R(S_i), I(B_j)) < U_j(P(S_j), R(S_j), I(B_j))$$

Equation 7 Partial Utility ordering for unsuccessful seller bidding

otherwise B_j could have selected S_i (if B_j was choosing first).

Since

$R(S_i) > R(S_j), I(B_i) > I(B_j)$, [Equation 2], [Equation 6], [Equation 7]

$$\Rightarrow P(S_i) > P(S_j)$$

Since this is true for any i, j such that $1 \leq i < j \leq t$ it means for any seller S_k , where $1 \leq k < t$, its closest competitor is S_{k+1} . Therefore, given $P(S_{k+1})$, the optimal price for S_k is the maximum (in order to maximize profit) price $P(S_k)$ for which if $i=k$ and $j=k+1$, then Equation 6 still holds. Indeed, if S_k offers slightly higher price $P(S_k)^+$, then S_{k+1} would undercut S_k by offering $P(S_{k+1})$, since we would have:

$$U_k(P(S_{k+1}), R(S_{k+1}), I(B_k)) > U_k(P(S_k)^+, R(S_k), I(B_k)).$$

Equation 8 Partial Utility ordering for unsuccessful seller bidding for consecutive pairs of sellers and buyers.

In the case of unemployment, the seller S_{t+1} will have to shout the minimal price (0.1 for our experiments) to maximize the probability of undercutting seller S_t . Therefore, the seller S_t will have to shout an $P(S_t)$ according to the assumption that $P(S_{t+1})=0.1$, and all the sellers will have to optimize the shouted prices according to the same assumption. In the case of full employment or overemployment, nobody can undercut seller S_t , therefore S_t will give the maximum price that meets the threshold of the minimum utility of the respective buyer B_t , so if there is no minimum utility requirement on the buyers' side, all sellers should be able ask for the maximum price. As it was discussed above, the minimum utility is the one enjoyed when the buyer transacts outside the online marketplace with a perfectly reputable sellers for the market price P_m .

This equilibrium state does not depend on the dynamics of the reputation algorithm itself. If the reputations of the sellers are stationary then the optimal pricing strategy would be the one that would price as close as possible to the optimal prices derived above at every trading period. However, since our system involves dynamically changing reputations, the dynamic pricing algorithms need to also adapt to these changes. Furthermore, we assume that sellers do not have complete information.

7.2 Optimal Strategy Sellers

We evaluate how socially optimal the different dynamic pricing strategies are by comparing their efficiency with the case where sellers have perfect information about the marketplace dynamics including: the numbers of sellers and buyers, the reputation of all the sellers, and the importance distributions of all the buyers. Although it is unrealistic that we would ever have such Optimal Strategy (OS) sellers, they provide a good benchmark for the evaluation of the intelligence and the social efficiency of our dynamic pricing algorithms.

The Optimal Strategy sellers would utilize all the information available to them in order to price according to the Nash Equilibrium described in the previous section. As it was explained above, the reputations of the sellers affect their overall profits mostly in unemployment environments, where only the most reputable will make transactions at equilibrium, rather than overemployment environments where all sellers will make transactions at equilibrium with reputation independent prices. Thus, this section and the following one will focus on unemployment environments.

In the case of unemployment, the Optimal Sellers will behave as follows:

All sellers know that the $t+1$ seller can try to undercut them by offering its services at minimal prices. Therefore, all sellers will have to match the utility offered by the $t+1^{st}$ seller when that seller bids $P_0 (=0.1)$. Consequently the optimal price for each seller would be such that:

$$U(B_j) = (1 - P_0)^{1-I(B_j)} R_{t+1}^{I(B_j)} = (1 - P(S_j))^{1-I(B_j)} R_j^{I(B_j)} \Rightarrow$$

$$\Rightarrow P(S_j) = 1 - \left[\frac{(1 - P_0)^{1-I(B_j)} R_{t+1}^{I(B_j)}}{R_j^{I(B_j)}} \right]^{\frac{1}{1-I(B_j)}}$$

Equation 9 Optimal Seller bidding with perfect information.

We now turn to some experiments comparing the optimal sellers with the Derivative and Reputation Followers.

7.3 Experimental comparison with Optimal Sellers

To observe closer the behavior of the agents we run experiments with few of them: we used 3 buyers and 10 sellers and we set the probability P of each buyer generating a task to be 1. This way we could easier track the matching of buyers and sellers. For simplicity, without loss of generality, we kept the importance sensitivities I of the buyers fixed: buyer 1 had I equal to 0.707, buyer 2 0.577, and buyer 3 0.5 - so the importance was decreasing with the buyer id. We first experimented with fixed reputations for the sellers, too. Their reputations were equal to their abilities, which were also a decreasing function of their id, shown at the two last columns of Table 1 and Table 2. All sellers started with prices equal to 0.1 - the minimum possible price they can charge. Table 1 and Table 2 show the equilibrium reached for derivative and Reputation Followers. The first column is the seller id (sellers 1 through 10), columns 2 and 3 show the average buyer id with whom each seller traded the first 50 iterations (-1 if the seller made no trades), and the total number of trades made by each seller during these iterations. For example in Table 1, seller 1 traded with "buyer" 1.5 (this is simply the arithmetic average of the buyers' ids with whom seller 1 trades), and had a total of 40 trades. Seller 6 traded only with buyer 3 a total of 13 trades. Similarly, columns 4 and 5 show the average buyer id and total number of trades during iterations 100-150, and columns 6 and 7 the same for iterations 750-800.

Table 1: Equilibrium reached by Derivative Followers when reputations and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations	Middle 50 iterations	Last 50 iterations	Seller Ability	Final Seller Reputation

	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	1.5	40	1	27	1	26	0.5	0.5
2	1.22	27	1.09	22	1.08	26	0.33	0.33
3	1.95	22	1.89	27	2	22	0.25	0.25
4	2.38	29	2.35	26	2.12	25	0.2	0.2
5	2.95	19	2.72	25	2.875	24	0.17	0.17
6	3	13	3	23	2.96	26	0.14	0.14
7	-1	0	-1	0	3	1	0.125	0.125
8	-1	0	-1	0	-1	0	0.11	0.11
9	-1	0	-1	0	-1	0	0.1	0.1
10	-1	0	-1	0	-1	0	0.09	0.09

Table 2 Equilibrium reached by Reputation Followers when reputations and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations		Middle 50 iterations		Last 50 iterations		Seller Ability	Final Seller Reputation
	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	1.32	44	1	24	1	26	0.5	0.5
2	2.16	49	1.19	32	1.11	27	0.33	0.33
3	2.11	38	2.08	37	2	24	0.25	0.25
4	2.94	18	2.74	38	2.04	24	0.2	0.2
5	3	1	3	19	3	25	0.17	0.17
6	-1	0	-1	0	3	24	0.14	0.14
7	-1	0	-1	0	-1	0	0.125	0.125
8	-1	0	-1	0	-1	0	0.11	0.11
9	-1	0	-1	0	-1	0	0.1	0.1
10	-1	0	-1	0	-1	0	0.09	0.09

According to the derivation above, the Optimal Sellers with complete information would trade as follows: seller 1 would trade with buyer 1, seller 2 with buyer 2, seller 3 with buyer 3, and sellers 4-10 would not trade. Instead of reaching this theoretical equilibrium, we notice that both Derivative Followers and Reputation Followers (Table 1 and Table 2 respectively) reach an "equilibrium" where sellers 1 and 2 "share" buyer 1 (half the times seller 1 trades with buyer 1 and seller 2 does no trade, and the other half the times seller 2 trades with buyer 1 and seller 1 does not trade), sellers 3 and 4 "share" buyer 2, sellers 5 and 6 "share" buyer 3, and sellers 7 through 10 do not trade. Notice that this equilibrium is not reached the first 50 iterations, and it is almost reached in 100 iterations (columns 4 and 5 are similar to columns 6 and 7, respectively). We also report the final prices charged by the sellers, to compare with the theoretically optimal ones given by Equation 9. Equation 9 gives that seller 1 should price 0.901479, seller 2 0.552088, seller 3 0.28, and sellers 4-10 0.1. Instead at the equilibrium reached the Derivative and Reputation

Followers charge similar prices as follows: seller 1: 0.919175, seller 2: 0.767465, seller 3: 0.542694, seller 4: 0.36816, seller 5: 0.229084, seller 6: 0.102303, and sellers 7 through 10: 0.1.

Table 3 Comparison of Final and Optimal prices with 3 buyers and 10 sellers when reputations and importance sensitivities are fixed.

Seller ID	Buyer Importance	Seller Abilities	Seller Final Prices	Optimal Prices
1	0.707107	0.5	0.919175	0.901479
2	0.57735	0.333333	0.767465	0.552088
3	0.5	0.25	0.542694	0.28
4		0.2	0.36816	0.1
5		0.166667	0.229084	-0.08
6		0.142857	0.102303	-0.26
7		0.125	0.1	-0.44
8		0.111111	0.1	-0.62
9		0.1	0.1	-0.8
10		0.090909	0.1	-0.98

We also experimented with dynamically changing reputations to study the equilibrium reached. The results are shown in Table 4 and Table 5. It is interesting to observe that the equilibrium reached by both types of sellers is the same as before, except that now the "rank" of the sellers is not based on their actual abilities, but on their reputations. Moreover, for Reputation Followers the final reputations of the sellers coincide with their true abilities, so the equilibrium reached is similar to that in Table 1 and Table 2. On the other hand, Derivative Followers reach different reputations and therefore different equilibrium: in Table 4 column 6 has many -1's mixed with normal ids, but the ids that are not -1 are still decreasing with the final sellers' reputations shown in the last column. In the general case of dynamically changing reputations, it is important that the dynamic pricing methods lead to equilibrium that not only agree with the theoretical one according to the reputations of the sellers, but also that the sellers' reputations coincide with their actual abilities.

Table 4 Equilibrium reached by Derivative Followers when reputations are dynamically changing and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations		Middle 50 iterations		Last 50 iterations		Seller Ability	Final Seller Reputation
	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	1.83	12	1.49	37	1	21	0.5	0.5
2	2.77	9	1.76	41	1	27	0.33	0.33
3	-1	0	-1	0	-1	0	0.25	0.01
4	-1	0	-1	0	-1	0	0.2	0.01
5	1.48	31	1.77	22	1.93	27	0.17	0.17
6	2.07	28	2.5	18	2.24	25	0.14	0.14
7	2.19	37	2.67	18	2.78	27	0.125	0.125

8	2.06	33	2.93	14	3	23	0.11	0.11
9	-1	0	-1	0	-1	0	0.1	0.01
10	-1	0	-1	0	-1	0	0.09	0.01

Table 5 Equilibrium reached by Reputation Followers when reputations are dynamically changing and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations		Middle 50 iterations		Last 50 iterations		Seller Ability	Final Seller Reputation
	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	-1	0	-1	0	1	24	0.5	0.50
2	1.94	50	1.72	50	1.07	27	0.33	0.33
3	-1	0	-1	0	1.96	23	0.25	0.25
4	2	50	2.14	50	2	20	0.2	0.2
5	-1	0	-1	0	2.63	16	0.17	0.17
6	-1	0	-1	0	3	16	0.14	0.14
7	-1	0	-1	0	3	13	0.125	0.125
8	-1	0	-1	0	3	5	0.11	0.11
9	2.06	50	2.14	50	3	6	0.1	0.1
10	-1	0	-1	0	-1	0	0.09	0.09

7.4 Profit Maximizing Reputation Followers

The results of the previous sections show that the Derivative Followers under-perform in cases of changing reputations, because the equilibrium prices do not match the seller's abilities. They are trapped in local maximum hills of their profit landscapes and they optimize their pricing for buyers of lesser quality sensitivities than the ones that match their abilities. On the other hand, reputation followers manage to have their reputations match their abilities, but they suffer from the same problem that both derivative followers and reputation followers have in the case of fixed reputations: the equilibrium reached is not the same as the theoretical one. Instead, the sellers oscillate their prices optimizing for two consecutive buyers rather than the one the same ranking of quality sensitivity (i.e. buyers 1, 2 and 3 buy from sellers (1,2), (3,4), and (5,6), instead of (1), (2), (3)). Therefore, we need a pricing mechanism that allows the sellers to escape from local maxima and learn the optimal prices for their abilities. For this purpose we designed dynamic pricing sellers that not only take into account their prices and reputations, but also their profits, and that compare prices, profits, and reputations over period of time so in a sense they have "memory" of the past. In particular the profit followers with memory behave as follows: for a given time window, say of length of 10 iterations, they measure their average prices, profit, and reputation over the most recent 10 iterations and of the previous 10 ones (from 20 iterations ago until 10 iterations ago). They then decide their next price based on the relative changes of their reputation, prices, and profits over these two periods. For example, if the profits, the prices, and the reputations increased relative to 20 to 10 iterations ago, the agents further increase their prices. If the profits decreased while the prices increased and the reputations decreased, they decrease their next price. We describe the actions taken under all market circumstances in table 6. For the cases that it is not clear whether to increase or decrease the price, the agents choose the average price of the past 2t iterations. We call these agents the Profit Maximizing Reputation Followers (PMRF). Table 6 shows the decision logic of the PMRF agents. In the four cases where the decision is ambiguous, the PMRF agents implement a divide and conquer approach, by choosing the mean of the

average price during the two consecutive periods. In future research we intend to weigh the changes of profits, prices and reputations and select a more optimal value than the mean, based on the incremental impact of prices and reputations on the buyers' perceived utilities. Such an approach requires that the sellers maintain a model for the utility functions of the buyers.

Table 6 Decision table for Profit Maximizing Reputation Followers

Profits	Price	Reputation	PMRF decision
Up	Up	Up	Up
Up	Up	Down	Down/up
Up	Down	Up	Down/up
Up	Down	Down	Down
Down	Up	Up	Down/up
Down	Up	Down	Down
Down	Down	Up	Down/up
Down	Down	Down	Down

We run simulations using profit followers with memory and we recorded the results shown in Table 7, Table 8, and Table 9. Table 7 shows the equilibrium results with stationary reputations, and Table 8 and Table 9 show the equilibrium results with changing reputations. As we can see from both Table 7 and Table 8 the sellers end up optimizing their pricing in order to trade with their respective buyers, in both the experiments with stationary, and non stationary reputations. Table 9 shows that the final seller prices are slightly higher than the optimal for all the sellers who are able to make transactions except the first one.

Table 7 Equilibrium reached by Profit Maximizing Reputation Followers when reputations and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations		Middle 50 iterations		Last 50 iterations		Seller Ability	Final Seller Reputation
	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	2.06	50	1.346939	49	1	48	1	0.5
2	1.638298	47	1.625	40	1.978261	46	2	0.33
3	1.911765	34	2.533333	30	2.851064	47	3	0.25
4	2.894737	19	3	31	3	9	4	0.2
5	-1	0	-1	0	-1	0	5	0.167
6	-1	0	-1	0	-1	0	6	0.14
7	-1	0	-1	0	-1	0	7	0.125
8	-1	0	-1	0	-1	0	8	0.11
9	-1	0	-1	0	-1	0	9	0.1
10	-1	0	-1	0	-1	0	10	0.09

Table 8 Equilibrium reached by Profit Maximizing Reputation Followers when reputations are dynamically changing and importance sensitivities are fixed. See text for explanation.

Seller ID	First 50 iterations		Middle 50 iterations		Last 50 iterations		Seller Ability	Final Seller Reputation
	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades	Average Buyer ID	Number of trades		
1	2.02	50	1.734694	49	1	48	1	0.5
2	1.54	50	1.787234	47	2	47	2	0.33
3	2.44	50	2.4	50	2.857143	42	3	0.25
4	-1	0	2.75	4	2.923077	13	4	0.2
5	-1	0	-1	0	-1	0	5	0.167
6	-1	0	-1	0	-1	0	6	0.14
7	-1	0	-1	0	-1	0	7	0.125
8	-1	0	-1	0	-1	0	8	0.11
9	-1	0	-1	0	-1	0	9	0.1
10	-1	0	-1	0	-1	0	10	0.09

Table 9 Comparison of the final prices of Profit Maximizing Reputation Followers and Optimal prices with 3 buyers and 10 sellers when reputations are changing and buyer sensitivities are fixed.

Seller ID	Buyer Importance	Seller Abilities	Seller Final Prices	Optimal Prices
1	0.707107	0.5	0.8504	0.901479
2	0.57735	0.333333	0.6507	0.552088
3	0.5	0.25	0.4697	0.28
4		0.2	0.3904	0.1
5		0.166667	0.1	-0.08
6		0.142857	0.1	-0.26
7		0.125	0.1	-0.44
8		0.111111	0.1	-0.62
9		0.1	0.1	-0.8
10		0.090909	0.1	-0.98

8 Conclusions

We studied dynamic pricing methods for an Agent mediated Marketplace in which users' reputations are established through collaborative mechanisms. We described the setup of our marketplace, and we studied the equilibria reached by simple Derivative Follower agents, which adjust their prices based on the success of the previous bid, and Reputation Follower agents who adjust their prices based on both their owners'

reputations and the success of their previous bid. We compared the performance of these sellers under different market conditions through simulations: in the case of unemployment, Reputation Followers do perform better than Derivative Followers since the first manage to undercut the second. However, in the case of overemployment where undercutting is not as important (since all agents get a contract as long as they price in a way that they are not "undercut" by the external market price P_m), Derivative Followers perform better than Reputation Followers do. It is therefore important for the agents to be able to identify the market conditions, according to the requests for bids they receive, and adapt their pricing strategies accordingly.

We also presented a theoretically "socially optimal" pricing equilibrium and compared it with the equilibrium reached by the derivative and Reputation Followers: when the reputations of the sellers are fixed, both Derivative and Reputation Followers manage to reach an equilibrium similar, but not the same, to the theoretical one. This is a finding similar to that of [Cliff, D. and Bruten, J. (1998)] in the sense that "zero intelligence plus" [Cliff, D. and Bruten, J. (1998)] is shown to lead to interesting market equilibria. In the case of dynamically changing reputations, Derivative Followers did not manage to reach such an equilibrium, while Reputation Followers did. Based on these experimental findings we designed a new dynamic pricing approach that lead to agents, which are Profit Maximizing Reputation Followers with memory. The experiments with these agents lead to equilibria that are close to the theoretically optimal ones. A future direction of research is to employ these agents to real marketplaces and to further study their behavior.

References

- [1] Chavez, A. and Maes, P. 1996. Kasbah: An Agent Marketplace for Buying and Selling Goods. Proceedings of the First International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology (PAAM'96). London, UK.
- [2] Cliff, D. and Bruten, J. 1998. More Than Zero Intelligence Needed for Continuous Double-Auction Trading, HP Technical Report HPL-97-157, In Proc. IFAC Symposium on Computation in Economics, Finance and Engineering (CEFES'98). University of Cambridge, England.
- [3] Friedman, E. and Resnick, P. 1998. The Social Cost of Cheap Pseudonyms: Fostering Cooperation on the Internet. In Proceedings of the 1998 Telecommunications Policy Research Conference. Alexandria, VA. Lawrence Erlbaum Associates, Inc., Mahwah, NJ.
- [4] Glickman, M. Parameter estimation in large dynamic paired comparison experiments. *In Journal of the Royal Statistical Society: Series C (Applied Statistics)* vol. 48, no. 3, pp 377-394, Blackwell Publishers, Oxford, UK.
- [5] Kephart, J., O., Hanson, J., E., and Greenwald, A., R. 2000. Dynamic Pricing by Software Agents. *In Computer Networks*, vol. 32 (6). Elsevier Science, Amsterdam, the Netherlands.
- [6] Kollock, P. 1999. The Production of Trust in Online Markets, *Advances in Group Processes*, vol. 16, E. J. Lawler, M. Macy, S. Thyne, and H. A. Walker (eds). JAI Press, Greenwich, CT.
- [7] Moukas, A., Zacharia, G., Guttman, R., and Maes, P. 2000. Agent-Mediated Electronic Commerce: An MIT Media Laboratory Perspective. *In International Journal of Electronic Commerce*, vol. 4 (3), M.E. Sharpe, Armonk, NY pp5-22.
- [8] Nash, J. 1950. The Bargaining Problem. *Econometrica* 18 pp155-62.
- [9] Oliver, J. 1997. Artificial agents learn policies for multi-issue negotiation. *International Journal of Electronic Commerce*, vol. 1 (4), M.E. Sharpe, Armonk, NY pp. 49-88.
- [10] Sairamesh J., and Kephart, J., O. 1998. Price Dynamics of Vertically Differentiated Information Markets. In Proceedings of *International Conference on Information and Computation Economics*, Charleston, SC.
- [11] Wang, D. 1999. Market Maker: an Agent-Mediated Marketplace Infrastructure, MEng Thesis, Massachusetts Institute of Technology, Cambridge, MA.
- [12] Zacharia, G., Evgeniou, T., Moukas, A., Boufounos P., and Maes P. 2000. Economics of Dynamic Pricing in a Reputation Brokered Agent Mediated Marketplace, Special Issue of the Electronic Commerce Research Journal on Intelligent Agents in Electronic Commerce, Liu, J., Ye, Y., (eds). Baltzer Science Publishers, Bussum, the Netherlands.
- [13] Zacharia, and Maes, P. 2000. Trust Management through Reputation Mechanisms, *In Trust in Agents, Applied Artificial Intelligence*, vol. 14 (9), Castelfranchi, C., Falcone, R., Sadighi, B., F., Tan, Y., H. (eds.) H. Taylor & Francis, London, UK.