

Reputation Based Buyer Strategy For Seller Selection For Both Frequent and Infrequent Purchases

Abstract. Previous research in the area of buyer strategies for choosing sellers in ecommerce markets has focused on frequent purchases. In this paper we present a reputation based buyer strategy for choosing sellers in a decentralized, open, uncertain, dynamic, and untrusted B2C ecommerce market for frequent and infrequent purchases. The buyer models the reputation of the seller after having purchased goods from it. The buyer has certain expectations of quality and the reputation of a seller reflects the seller's ability to provide the product at the buyer's expectation level, and its price compared to its competitors in the market. The reputation of the sellers and the price quoted by the sellers are used to choose a seller to transact with. We compare the performance of our model with other strategies that have been proposed for this kind of market. Our results indicate that a buyer using our model experiences a slight improvement for frequent purchases and significant improvement for infrequent purchases.

Keywords: autonomous agents, intelligent agents, ecommerce, learning, reputation

1 Introduction

Our work considers decentralized, open, dynamic, uncertain and untrusted market places. Sellers sell products and the quality and the price of product varies across sellers. The goal for the buyer is to purchase a product from a seller who meets its expectations of quality and service and purchase it at the lowest price possible in the market. At the same time the buyer wants to reduce its chances of interacting with dishonest and poor quality sellers. In an open market, sellers and buyers can enter and leave the market anytime. In a dynamic market the players in the market need not exhibit the same behavior all the time; sellers can vary the price and the quality in various transactions. Untrusted market implies there could be dishonest sellers in the market. By uncertain market we mean that the buyers can gauge the quality of the product after actually receiving the product. There could be a onetime transaction between the buyer and seller or multiple transactions between them. There is no limitation on the number of sellers and the buyers in the market. These characteristics are typical of a traditional commerce market and hence we consider a similar environment for our electronic market.

It is not possible to preprogram an agent to operate under these conditions, or to know beforehand who the best seller for a buyer is, as new sellers are entering the market, the lowest priced seller may not necessarily be the best seller, and sellers could be lying. Agents have to be equipped with abilities to make the most rational decision based on all the information that they can gather. They should be able learn from their past experiences.

Recent research has developed intelligent agents for ecommerce applications [1], [2], [3], [4], [5], [7], [8], [9]. However, as Tran [7] summarizes, the agents in [4], [5] are not autonomous, agents in [1], [2], [3], [4] do not have learning abilities, the agents in [9] have significant computational costs, and the agents in [1], [2], [3], [4],

[5], [9] do not have the ability to deal with deceptive agents. Tran and Cohen's, and Tran's work [7], [8] addressed these shortcomings by developing a strategy for buying agents using reinforcement learning and reputation modeling of the sellers. However their model builds reputation slowly and the buyer has to interact with a seller several times before the seller is considered reputable. This model works well where the buyer has to make repeated transactions with the sellers during frequent purchases. The performance of this model deteriorates for infrequent purchases as the buyer has to purchase several times from a seller before making its decision about the seller. When the buyer is purchasing a product on an infrequent basis it needs to quickly identify reputed sellers.

We present reputation based modeling of a seller by the buyer which can work for frequent as well as infrequent purchases in a B2C ecommerce market. We compared the performance of buying agents using our model, reinforcement learning[9] and reputation based reinforcement learning [7], [8]. Our results show that buying agents using our model improved their performance slightly for frequent purchases and showed a significant improvement for infrequent purchases, making our approach better suitable for all kinds of buyers.

2 Methodology

We consider decentralized, open, dynamic, uncertain and untrusted market places. Buyer agents model the sellers' reputation based on their direct interactions with them. The buyer has certain expectations of quality and the reputation of a seller reflects the seller's ability to provide the product at the buyer's expectation level, and its price compared to its competitors in the market. The buyer's goal is to purchase from a seller who will maximize its valuation of the product, which is a function of the price and quality of the product. At the same time it wants to avoid interaction with dishonest or poor quality sellers in the market. The reputation of the seller is used to weed out dishonest or poor quality sellers.

In this paper we use the following notation: Subscript represents the agent computing the rating. Superscript represents the agent about whom the rating is being computed. The information in the parenthesis in the superscript is the kind of rating being computed. For example, every time buyer b purchases a product from the seller, it computes a direct trust (di) rating $T_b^{s(di)}$ of the seller s by buyer b . The trust rating of seller s by buyer b is computed as shown in equation 1.

$$T_b^{s(di)} = \begin{cases} \frac{q_{act}}{q_{exp}} - \left(\frac{p_{act} - p_{avg}}{p_{max}} \right) & \text{if } q_{act} \geq q_{min} \text{ and } p_{act} \geq p_{avg} & (a) \\ \frac{q_{act}}{q_{exp}} & \text{if } q_{act} \geq q_{min} \text{ and } p_{act} < p_{avg} & (b). \\ \frac{q_{act}}{q_{exp}} - \left(\frac{p_{act} - p_{min}}{p_{max} - p_{min}} \right) & \text{if } q_{act} < q_{min} & (c) \end{cases} \quad (1)$$

where q_{act} is the actual quality of the product delivered by seller s , q_{exp} is the desired expected quality and q_{min} is the minimum quality expected by b . p_{act} is the price paid by b to purchase the product from seller s . p_{min} is the minimum price quote, p_{max} is the maximum price quote received and p_{avg} is the average of the price quotes received by the buyer for this product.

The trust rating should be proportional to the degree the quality delivered by the seller meets the buyer's expectations and the price paid to purchase the product. If there are two sellers, $s1$ and $s2$, who can meet the buyer's expectation for the quality of the product, and $s1$'s price is lower than $s2$, then $s1$ should get a higher rating than $s2$. Similar to [5] and [6] we make the common assumption that it costs more to produce a higher quality product. So when considering the price charged by a seller, if the seller meets the buyer's minimum expectation for quality, and if the price is greater than the average price quoted, then the difference between the seller's price and the average price quoted is weighed against the maximum price quoted for that product (part (a) of the equation). On the other hand if the price of the seller is below the average price (which can happen if the other sellers are trying to maximize their profits or there are too many low quality sellers) then the rating for this seller is computed based on its quality alone (part (b) of the equation). If the seller's quality does not meet the buyer's expectation then the difference of seller's price and the minimum price quoted is compared to the difference between the maximum and the minimum price quoted to penalize the seller more severely (part (c) of the equation).

This model makes the assumption that buyer b expects the highest quality and in the best case q_{act} can be equal to q_{exp} and it costs more to produce higher quality products. From the above equations it can be seen that $T_b^{s(di)}$ ranges from $[-1, 1]$. In the best case, b gets the expected quality at the lowest price and $T_b^{s(dimax)} = 1$. In the worst case $q_{act} = 0$ and b pays the maximum price quoted and $T_b^{s(dimin)} = -1$.

If the buyer has not interacted with the seller then $T_b^{s(di)} = 0$ for that seller and such a seller is referred to as a new seller.

Whenever buyer b is evaluating a list of sellers for purchase decisions it computes $T_b^{s(diavg)}$, the average rating for each seller s from its past interactions. $T_b^{s(diavg)}$ is computed as the weighted mean of its past n recent interactions.

$$T_b^{s(diavg)} = \frac{1}{W} \sum_{i=1}^n w_i T_b^{s(di)} \quad (2)$$

Where

$$w_i = \frac{t_{cur}}{t_{cur} - t_i} \quad (3)$$

$$W = \sum_{i=1}^n w_i \quad (4)$$

Where $T_{b(i)}^{s(di)}$ is the rating computed for a direct interaction using equation 1. Subscript i in parenthesis indicates the i^{th} interaction. w_i is the importance of the rating

in computing the average. Recent ratings should have more importance. Hence the weight of a rating is inversely proportional to the difference between the time a transaction happened t_i to the current time t_{cur} .

The buyer has threshold values θ and ω for the direct trust ratings to indicate its satisfaction or dissatisfaction with the seller respectively. θ and ω are set by the buyer and $\theta > \omega$ and θ and ω are in the range $[-1, 1]$. Buyer chooses sellers whose average direct trust rating is greater than or equal to θ and considers them to be reputable, does not choose sellers whose average direct trust rating is less than or equal to ω and considers them to be disreputable. It is unsure about sellers whose average direct trust ratings are between ω and θ and will consider them again only if there are no reputable or new sellers to consider. From the list of sellers who have submitted price bids, reputable sellers whose $T_b^{s(diavg)}$ is above the satisfaction threshold θ are identified as potential sellers. Buyer includes new sellers into the list of potential sellers to be able to quickly identify a good seller.

The buyer's valuation function for the product is a function of the price a seller is currently quoting and the quality that has been delivered in the past. For a seller with whom the buyer has interacted before, the quality is the average of the quality delivered in the past interactions. For a seller with whom the buyer has not interacted directly, the quality is set to the expected quality. From the list of potential sellers, buyer chooses a seller who maximizes its product valuation function.

3 Related Work

We compare our model to [7], [8], [9] as their and our work consider a similar market environment with autonomous buying agents who learn to identify sellers to transact with. [9] use reinforcement learning strategy and [7], [8] use reinforcement learning with reputation modeling of sellers. Our model provides a different method of computing reputation and does not use reinforcement learning strategy.

Vidal and Durfee's [9] economic model consists of seller and buyer agents. The buyer has a valuation function for each good it wishes to buy which is a function of the price and quality. The buyer's goal is to maximize its value for the transaction. Agents are divided into different classes based on their modeling capabilities. 0-level agents base their actions on inputs and rewards received, and are not aware that other agents are out there. 1-level agents are aware that there are other agents out there, and they make their predictions based on the previous actions of other agents. 2-level agents model the beliefs and intentions of other agents. 0-level agents use reinforcement learning. The buyer has a function f for each good that returns the value that the buyer expects to get by purchasing the good at price p . This expected value function is learned using reinforcement learning as $f = f + \alpha(v - f)$ where α is the learning rate, initially set to 1 and reduced slowly to minimum value. The buyer picks a seller that maximizes its expected value function f . Our market model is extended into a more general one by having sellers offer different qualities and by the existence of dishonest sellers in the market. The buyers use the reputation of the sellers to avoid dishonest sellers and reduce their risks of purchasing low quality goods. The reputation of the sellers is learned based on direct interactions.

Tran and Tran and Cohen develop learning algorithms for buying and selling agents in an open, dynamic, uncertain and untrusted economic market [7], [8]. They use Vidal and Durfee's [9] 0-level buying and selling agents. The buying and selling agents use reinforcement learning to maximize their utilities. They enhance the buying agents with reputation modeling capabilities, where buyers model the reputation of the sellers. The reputation value varies from -1 to 1. A seller is considered reputable if the reputation is above a threshold value. The seller is considered disreputable if the reputation value falls below another threshold value. Sellers with reputation values in between the two thresholds are neither reputable nor disreputable. The buyer chooses to purchase from a seller from the list of reputable sellers. If no reputable sellers are available, then a seller from the list of non disreputable sellers is chosen. Initially a seller's reputation is set to 0. The seller's reputation is updated based on whether the seller meets the demanded product value. If the seller meets or exceeds the demanded product value then the seller is considered cooperative and its reputation is incremented. If the seller fails to meet the demanded product value then the seller is considered uncooperative and its reputation is decremented. This model builds reputation slowly. So the buyer has to interact with a seller several times before the reputation of the seller crosses the threshold value. This model works well where the buyer has to make repeated transactions with the sellers, but a buyer cannot utilize this model when making infrequent purchases.

4 Experiments and Results

For our experiments we developed a multi-agent based simulation of an electronic market with autonomous buying, selling agents, and a matchmaker. Sellers upon entering the market register with a matchmaker [6] regarding the products that they can supply. When a buyer wants to purchase a product, it obtains a registered list of sellers selling this product from the matchmaker and sends a message to each of the sellers in the list to submit their bids for the product p . Sellers who are interested in getting the contract submit a bid which includes the price. The buyer waits for a certain amount of time for responses and then evaluates the bids received to choose a seller to purchase from.

The following parameters were set. Quality q sold across the sellers ranges from [10, 50] and varies in units of 1. Buyer expects a minimum quality of $40(q_{min}=40)$. The price of a product for an honest seller is $pr = q \pm 10\%q$. Like Tran [7] we make the assumption that it costs more to produce high quality goods. We also make the reasonable assumption that the seller may offer a discount to attract the buyers in the market or raise its price slightly to increase its profits. Hence the price of the product is set to be in the range of 90% -110% of the quality for an honest buyer. A dishonest buyer on the other hand may charge higher prices. The buyer's valuation of the product is a function of the quality and the price and for our simulation we set it as $3 * quality - price$. The buyer's valuation function reflects the gain, a buyer makes from having purchased a product from a seller. Each time a buyer purchases a product from a seller its product valuation is computed and we consider this as the buyer's gain for having purchased from that seller.

We compared the performances of four buyers:

1. *F&NF (Frequent and Infrequent) Buyer*: - This buying agent uses the buying strategy as described in our model. Buyer's desired expected quality is $q_{exp} = 50$. Acceptable quality for a buyer is from [40, 50]. Non acceptable quality is from [10-39]. Maximum price p_{max} quoted by honest seller would be 55 and minimum price p_{min} quoted would be 9. The average price p_{avg} would be 32. Threshold values θ for a seller to be considered reputable and ω for a seller to be considered disreputable values can be computed as:

The buyer is expecting at least a quality of 40. In the worst case it can get this at the highest price that can be charged by a honest seller which would be 44. From equation 1(a) the trust rating for that seller would be

$$\frac{40}{50} - \left(\frac{44 - 32}{55} \right) = 0.581. \quad (5)$$

So we set $\theta = 0.58$. For new sellers the trust rating is set to 0. These buyers should not come under the category of disreputable sellers. So we set the threshold value for a seller to be considered unacceptable as -0.1. So $\omega = -0.1$

2. *Tran Buyer*: - This buying agent uses the buying strategy as described in Tran and Cohen [8]. The threshold for seller to be considered reputable is set to 0.5 and for seller to be considered disreputable is set to -0.9 as described in their work.
3. *RL Buyer*: - This buying agent uses a Reinforcement learning strategy as described for 0-level buying agent in Vidal and Durfee [9].
4. *Random Buyer*: - This buying agent chooses a buyer randomly.

We populated the market with 12 sellers belonging to one of the six categories with the price and quality properties as shown (two agents per category):

1. *Honest Acceptable (HA)*: - Each seller offers a quality in the range [40-50]. Price is between 90-110% of the quality they are selling.
2. *Honest Not Acceptable (HNA)*: - Each seller offers a quality in the range [10-39]. Their price is between 90 -110% of the quality they are selling.
3. *Overpriced Acceptable (OPA)*: - Each seller offers a quality in the range [40-50]. Price is between 111-200% of the quality they are selling.
4. *Overpriced Not Acceptable (OPNA)*: - Each seller offers a quality in the range [10-39]. Their price is between 111-200% of the quality they are selling.
5. *Inconsistent*: - Each seller offers a quality in the range [10-50]. Price is between 90-110% of the quality they are selling.
6. *Dishonest*: - This category of sellers in their first sale to a buyer offer acceptable quality q [40-50] charging a price $pr = q \pm 10\%q$. In their subsequent sales to that buyer they reduce the quality q to be in the range [10-25]. However their price still remains high. Price $pr = q1 \pm 10\%q1$ where $q1$ is in the range [40 -50].

The data from the experiments was collected over 100 simulations. In each simulation, each buying agent conducted 500 transactions. In each transaction they purchased product p by querying the seller list from the matchmaker, obtain price quotes from different sellers and utilize their buying strategy to choose a seller. We compared the performances of the various buying agents on the following parameters.

- How long it took them to learn to identify high quality low priced sellers. We want the buying agents to identify high quality sellers offering low prices as soon as

possible. If the buyer is able to identify high quality sellers quickly then the same strategy can be used when making infrequent purchases.

- The average gain as the number of purchases of product p is increased. If the average is consistently high means that the buyer is interacting with high quality sellers offering low prices most often. If the average gain is high earlier on implies that the buyer has identified high quality low price sellers quickly.

Figures 1-3 show the gain versus transactions for each type of buyer (because of space considerations we are not showing the plot of the gain vs. a random buyer, since the gain simply constantly fluctuates):




Fig. 1. Gain Vs Transaction for a F&NF Buyer




Fig. 2. Gain Vs Transaction for Tran Buyer




Fig. 3. Gain Vs Transaction for RL Buyer

Table 1 shows the number of purchases made by a buyer from each seller type.

Table 1. Buyer seller interaction

	HA	HNA	OPA	OPNA	INC	DIS
Rsk Buyer	488	2	2	2	2	4
Tran Buyer	451	7	23	5	8	6
RL Buyer	420	16	15	13	17	16
Random Buyer	86	88	82	83	69	92

Acceptable quality sellers can offer qualities anywhere between 40-50. The lowest gain from purchasing from a honest seller offering at the lowest end of good quality range and charging its highest price is 76 ($3 \cdot 40 - 44$). When the gain from purchasing from a seller is 76 and above, it means the buyer is purchasing from a high quality low priced seller. From figures 1-3 it can be seen that F&NF Buyer, Tran Buyer and RL Buyer learn although at different rates to identify high quality low priced sellers. After having learned, they consistently interact with high quality low priced sellers. This is confirmed by the fact that highest number of purchases are made from honest acceptable sellers as shown in table 1. Random Buyers never learn and that is to be expected as they are choosing sellers randomly. F&NF Buyer learns to identify high quality low priced sellers very quickly in about 15 transactions or purchases. Tran Buyers take about 60 transactions to learn and RL Buyer learns in about 250 transactions. If the buyers were to purchase the product infrequently then the F&NF Buyer strategy would work better than the RL Buyer or Tran Buyer strategy as it requires the least number of transactions to learn.

Figure 5 shows the average gain versus the number of purchases for different buyers.




Fig. 5. Average Gain versus Number of Purchases for different buyers

In the beginning, average gains are fluctuating as the buyers employing a non-random strategy are learning and Random Buyer is choosing sellers randomly. F&NF Buyer is the quickest to learn and its average gain raises sharply earlier on compared to the other two learning agents. As RL Buyer takes a long time to learn, its average gain at the end is still lower than the F&NF or Tran Buyer. Since Random Buyer purchases randomly from various types of sellers, its average is consistently the lowest. In the first half of the Figure 5 it can be seen that when the purchases are fewer, the average gain for the F&NF Buyer, once its learning phase is completed, is higher than the other buying agents. So, if the buyers were to purchase the product infrequently, then the F&NF Buyer strategy works better than the RL or Tran Buyer strategy. As the number of purchases increases, F&NF Buyer still has the highest average gain with the Tran Buyer's average gain coming very close to it at very high number of purchases.

5 Conclusions and Future Work

We presented a model for a buyer to maintain the seller reputation and strategy for buyers to choose sellers in a decentralized, open, dynamic, uncertain and untrusted multi-agent based electronic markets. The buyer agent computes a seller agent's reputation based on its ability to meet its expectations of product, service, quality and price as compared to its competitors. We show that a buying agent utilizing our model of maintaining seller reputation and buying strategy does better than buying agents employing strategies proposed previously for frequent as well as for infrequent purchases. For future work we are looking at how the performance of buying agent can be improved for extremely infrequent purchases.

References

1. A. Chavez, and P. Maes, "Kasbah: An Agent Marketplace for Buying and Selling Goods." In Proceedings of the 1st Int. Conf. on the Practical Application of Intelligent Agents and Multi-Agent Technology, London, 1996.
2. A Chavez, D. Dreilinger, R. Guttman. And P. Maes, " A real-Life Experiment in Creating an Agent Marketplace", In Proceedings of the Second International Conference on the Practical Application of Intelligent Agents and Multi-Agent Technology, 1997.
3. C. Goldman, S. Kraus and O. Shehory, "Equilibria strategies: for selecting sellers and satisfying buyers," Lecture Notes in Artificial Intelligence, Vol. 2182, M. Klusch and F. Zambonelli (Eds.), Springer, 166-177, 2001.
4. R.B Doorenbos, Etzioni, and D. Weld, "A Scalable Comparison-Shopping Agent for the World Wide Web", In Proceedings of the First International Conference on Autonomous Agents, 39-48, 1997.
5. B. Krulwich, " The bargainer agent: comparison price shopping on the Internet", in J. Williams, editor, Bots, and other Internet Beasties, 257-263, Macmillan Computer Publishing, 1996.
6. D. Kuokka and L. Harada, "Supporting Information Retrieval via Matchmaking," Working Notes of the AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments, 1995.
7. T. Tran, Reputation-oriented Reinforcement Learning Strategies for Economically-motivated Agents in Electronic Market Environment, Ph.D. Dissertation, University of Waterloo, 2003.
8. T. Tran and R. Cohen, "Improving User Satisfaction in Agent-Based Electronic Marketplaces by Reputation Modeling and Adjustable Product Quality," Proc. of the Third Int. Joint Conf. on Autonomous Agents and Multi Agent Systems (AAMAS-04), Volume 2, 828-835, 2004.
9. J.M. Vidal and E.H. Durfee, "The impact of Nested Agent Models in an Information Economy", Proc. of the Second Int. Conf. on Multi-Agent Systems, 377-384, 1996.

Formatted: Bullets and Numbering