

Simulating the Effect of Reputation Systems on E-markets

Audun Jøsang, Shane Hird, and Eric Faccer

Distributed Systems Technology Centre*

Queensland University of Technology, GPO Box 2434, Brisbane Qld 4001, Australia
{ajosang, shird, efaccer}@dstc.edu.au

Abstract. Studies show that reputation systems have the potential to improve market quality. In this paper we report the results of simulating a market of trading agents that uses the beta reputation system for collecting feedback and computing agents' reputations. The simulation confirms the hypothesis that the presence of the reputation system improves the quality of the market. Among other things it also shows that a market with limited duration rather than infinite longevity of transaction feedback provides the best conditions under which agents can adapt to each others change in behaviour.

1 Introduction

Individuals and organisations take risks by participating in e-commerce. Transactions are often conducted with people or organisations who are strangers to each other and who often have an unknown record of past behaviour. This situation provides plenty of opportunities for irresponsible and fraudulent players to provide low quality or deceitful services resulting in a large number of complaints about e-commerce players being reported [5]. A study by Consumers International [3] involved researchers from 11 countries placing 151 orders at sites in 17 countries. They found that 9% of goods ordered never arrived. In 20% of cases the amount actually charged was higher than expected, and there were problems obtaining a refund in 21% of purchases, despite the fact that the sites in question advertised that refunds were available. When little or no information about potential transaction partners can be obtained honest players are left in the dark when making decisions about e-business transactions. Hence, there is a need to design and introduce systems for recording the behaviour of e-market participants and make it available to potential transaction partners for decision making purposes. This is in essence the role of reputation systems.

A reputation system gathers and aggregates feedback about participants behaviour in order to derive measures of reputation. In the centralised approach which we will use in our market simulation study, feedback is received and stored by a feedback collection centre. Fig.1 below shows a typical reputation framework where X and Y denote transaction partners.

* The work reported in this paper has been funded in part by the Co-operative Research Centre for Enterprise Distributed Systems Technology (DSTC) through the Australian Federal Government's CRC Programme (Department of Industry, Science & Resources).

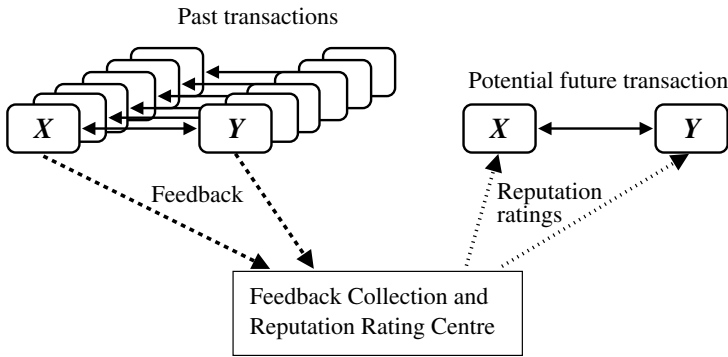


Fig. 1. General framework for a centralise reputation system

After a transaction is completed, the agents provide feedback about each other's performance during the transaction. The reputation centre collects feedback from all the agents and continuously updates each agent's reputation rating as a function of the received feedback. Updated reputation ratings are provided online for all the agents to see, and can be used by the agents to decide whether or not to transact with a particular agent. The two fundamental aspects of reputation systems are:

1. *Propagation mechanism* that allow participants to provide feedback about transaction partners, as well as to obtain reputation measures of potential transaction partners.
2. *Reputation engine* that derive measures of reputation for each participant based on transaction feedback and other information.

Studies show that reputation systems can improve the quality of e-markets, and that the reputation of a particular agent can have positive effects on the agent's gain as a result of participating in the marketplace. The quality of the market can for example be defined in terms of efficiency and honesty as described by Bolton *et al.* (2002) [1].

The quality of the market and the gain of individual participants as described above represent very specific and easily measurable properties. Reputation systems thus address both the overall welfare of communities and the prosperity of individual participants. To stimulate environments where reputation can develop thus seems to be a promising way of creating healthy markets. The purpose of this study is to analyse the effect of using a reputation system in a simulated e-market.

2 The Reputation System

This simulation study was done using the Beta Reputation System [4] which provides a flexible framework for integrating reputation services into e-commerce applications. The reputation rating described in Section 2.2 is based on the beta probability distribution function which will be briefly described next.

2.1 The Beta Distribution

The beta distribution function can be used to derive probability distributions based on observations of past events. The mathematical analysis leading to this result can be found in many text books on probability theory, e.g. Casella & Berger 1990[2] p.298, and we will only present the results here.

The beta-family of distributions is a continuous family of distribution functions indexed by the two parameters α and β . The beta distribution $f(p; \alpha, \beta)$ can be expressed using the gamma function Γ as:

$$f(p; \alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} p^{\alpha-1} (1-p)^{\beta-1} \quad \text{where } 0 \leq p \leq 1, \alpha > 0, \beta > 0 \quad (1)$$

with the restriction that the probability variable $p \in [0, 1]$ if $\alpha < 1$, and $p \in [1, 0]$ if $\beta < 1$. The probability expectation value of the beta distribution is given by:

$$E(p) = \frac{\alpha}{\alpha + \beta} \quad (2)$$

Let us consider a process with two possible outcomes $\{x, \bar{x}\}$ and let r be the observed number of outcome x and let s be the observed number of outcome \bar{x} . Then the probability distribution of observing outcome x in the future can be expressed as a function of past observations by setting:

$$\alpha = r + 1 \quad \text{and} \quad \beta = s + 1 \quad \text{where } r, s \in \mathbb{N}_0 \quad (3)$$

As an example, a process with two possible outcomes $\{x, \bar{x}\}$ that has produced 7 x and 1 \bar{x} will have a beta distribution expressed as $f(p; 8, 2)$ which is plotted in Fig.2.




Fig. 2. Beta function of event x after 7 observations of x and 1 observation of \bar{x} .

This curve expresses the uncertain probability that the process will produce outcome x during future observations. The probability expectation value according to Eq.(2) is $E(p) = 0.8$. This can be interpreted as saying that the relative frequency of outcome x in the future is somewhat uncertain, and that the most likely value is 0.8.

The variable p is a probability variable, so that for a given p the probability density $f(p \in \mathbb{P})$ represents second order probability. The first-order variable p represents the probability of an event, whereas the density $f(p \in \mathbb{P})$ represents the probability that the first-order variable has a specific value. Since the first-order variable p is continuous, the second-order probability $f(p \in \mathbb{P})$ for any given value of $p \in [0, 1]$ is vanishingly small and therefore meaningless as such. It is only meaningful to compute $\int_{p_1}^{p_2} f(p \in \mathbb{P})$ for a given interval $[p_1, p_2]$, or simply to compute the expectation value of p . Below we will define a reputation rating that is based on the expectation value. This provides a sound mathematical basis for combining feedback and for expressing reputation ratings.

2.2 The Reputation Rating

The concept of probability distribution is ideal for mathematical manipulation, but less so for communicating a reputation rating to human users. A more simple representation than the probability distribution is therefore needed. Most people are familiar with the notion of a probability value, and the probability expectation value $E(p)$ from Eq.(2) therefore seems very suitable. This would give a reputation rating in the range $[0, 1]$ and this will be used in the simulations. It should be noted however that the reputation rating can be scaled to any range as mentioned in [4].

The reputation rating will depend on the feedback as well as a base rate that will be equal for all participants in a market. This will be implemented using base rate parameters denoted by (r_{base}, s_{base}) . This will replace the fundamental base rate of $1/2$ of the beta probability distribution. The base rate and the feedback parameters will together determine the reputation rating according to the expression below.

Definition 1 (Reputation Rating). Let $(r(Z), s(Z))$ represent target Z 's reputation parameters, and (r_{base}, s_{base}) represent the base rate parameters. Then the function $Rep(Z)$ defined by:

$$Rep(Z) = \frac{r(Z)+r_{base}}{r(Z)+s(Z)+r_{base}+s_{base}} \tag{4}$$

is called Z 's reputation rating.

The reputation rating can be interpreted as a measure of reputation, that may be used as an indication of how a particular agent is expected to behave in future transactions. The parameter $r(Z)$ represents the amount of positive feedback about Z , and the parameters $s(Z)$ represents the amount of negative feedback about Z .

2.3 Combining Feedback

Reputation systems must be able to combine feedback from multiple sources. In the beta reputation system this can be done by simply accumulating all the received r and s parameters from the feedback providers. Assume two agents X and Y giving feedback about the same target Z . Target Z 's reputation parameters by X and by Y can then be expressed as $(r^X(Z), s^X(Z))$ and $(r^Y(Z), s^Y(Z))$ respectively. By simply adding the reputation parameters for Z the combined reputation parameters emerge.

Definition 2 (Combining Feedback).

Let $(r^X(Z), s^X(Z))$ and $(r^Y(Z), s^Y(Z))$ be two different sets of reputation parameters on Z resulting from X and Y 's feedback respectively. The reputation parameters $(r^{X \circ Y}(Z), s^{X \circ Y}(Z))$ defined by:

$$\begin{aligned} 1 \triangleright r^{X \circ Y}(Z) &= r^X(Z) + r^Y(Z) \\ 2 \triangleright s^{X \circ Y}(Z) &= s^X(Z) + s^Y(Z) \end{aligned} \tag{5}$$

is then called Z 's combined reputation parameters by X and Y .

2.4 Providing and Collecting Feedback

After each transaction, a single agent can provide both positive and negative feedback simultaneously, in the form of the parameters $r \in [0, 1]$ and $s \in [0, 1]$ respectively. The purpose of providing feedback as a pair (r, s) is to reflect the idea that an agent's performance in a transaction can, for example, be partly satisfactory. This could be expressed through a feedback with $(r, s) = (1, 0)$ which carries the weight of 1, whereas a feedback with $(r, s) = (0, 0)$ is weight-less, and would be as if no feedback was provided at all. In general the sum $r + s$ can be interpreted as the weight of the feedback.

In our simulation the buyers will provide the feedback $(r, s) = (1, 0)$ in case the item is shipped, and the feedback $(r, s) = (0, 1)$ in case the item is not shipped, so that the weight will always be 1.

2.5 Feedback Longevity

Old feedback may not always be relevant for the actual reputation rating, because the agent may change its behaviour over time. What is needed is a model in which old feedback is gradually forgotten or diminished. This can be achieved by introducing a feedback longevity factor $\alpha \in [0, 1]$ which can be adjusted in order to determine how fast feedback will diminish. The feedback longevity principle is very similar to the feedback forgetting principle described in [4]. The idea is that neither a feedback longevity $\alpha = 0$ (where all feedback is immediately forgotten) nor a feedback longevity $\alpha = 1$ (where all feedback is remembered forever) will result in a healthy market. Rather, the longevity factor must have a value which allows market players to, for example, repair a damaged reputation by improving their behaviour.

Let Z 's reputation parameters resulting from feedback received after i transactions be denoted by $(r_i(Z), s_i(Z))$. The feedback longevity factor $\alpha \in [0, 1]$ can be introduced so that when adding feedback from transaction $(i + 1)$ denoted by $(r(Z), s(Z))$ the new updated reputation parameters can be computed as:

$$\begin{aligned} 1 \triangleright r_{i+1}(Z) &= (\alpha r_i(Z) + (1 - \alpha)(r(Z))) \leq 1 \\ 2 \triangleright s_{i+1}(Z) &= (\alpha s_i(Z) + (1 - \alpha)s(Z)) \leq 1 \end{aligned} \quad \text{where } 0 \leq \alpha \leq 1 \tag{6}$$

This is a recursive algorithm which makes it possible to keep the reputation parameters of all agents without having to store every single feedback.

3 Description of the Market Simulation

3.1 Computing Transaction Gains

The market consists of 3 sellers and 10 buyers who trade the same commodity item. The market activity is structured into sessions of 10 rounds each. In each round every buyer will attempt to buy exactly one item. A transaction only happens when a buyer finds a seller satisfying his buying criteria. A buyer who receives the item provides a positive feedback $(r(S) \ominus s(S)) = (1 \ominus 0)$, and a buyer who does not receive it provides negative feedback $(r(S) \ominus s(S)) = (0 \ominus 1)$. The seller's reputation is then updated using Eq.(6), and the change becomes visible to the market at the beginning of the next round.

At the start of each session each seller S defines the selling price $\text{Pr}c(S)$ and honesty level $\text{Hst}(S)$. Every seller S will also have a reputation rating $\text{Rep}(S)$. Items are shipped with probability $\text{Hst}(S)$, and not shipped with probability $(1 \otimes \text{Hst}(S))$.

The production cost per item is $\text{Cst}_{\text{prod}} = \80 . In case the item is shipped the actual transaction gain of seller S can be computed as:

$$g^{\text{T}}(S) = \text{Pr}c(S) \otimes \text{Cst}_{\text{prod}} \triangleright \quad (7)$$

The seller will thus make a positive gain $g^{\text{T}}(S)$ from a transaction as long as $\text{Pr}c(S) > \text{Cst}_{\text{prod}}$. By defining a selling price the seller is in fact defining the gain per transaction because the production cost is fixed. In case the seller does not ship the item, the actual transaction gain is simply the selling price according to:

$$g^{\text{T}}(S) = \text{Pr}c(S) \triangleright \quad (8)$$

The intrinsic value of the item to every buyer is denoted by $\text{Val} = \$260$. The buyers are faced with the risk that the seller does not ship, and the buyer's expected transaction gain when dealing with a specific seller S is denoted by $g_{\text{Exp}}^{\text{T}}(S)$ which can be computed as:

$$\begin{aligned} g_{\text{Exp}}^{\text{T}}(S) &= (\text{Val} \otimes \text{Pr}c(S)) \triangleleft \text{Rep}(S) \otimes \text{Pr}c(S) \triangleleft (1 \otimes \text{Rep}(S)) \\ &= \text{Val} \triangleleft \text{Rep}(S) \otimes \text{Pr}c(S) \triangleright \end{aligned} \quad (9)$$

At the start of each session all the buyers will define the *risk aversion* denoted by $\text{Avs}(B) \sqcap [0 \ominus 1]$ where $\text{Avs}(B) = 0$ means that the buyer is happy to take any risk, and $\text{Avs}(B) = 1$ means that the buyer does not want to take any risk at all. By setting a risk aversion the buyer B requires that the seller S has a reputation rating satisfying:

$$\text{Rep}(S) \sqcap \text{Avs}(B) \triangleright \quad (10)$$

A buyer will thus choose the seller who yields the highest expected gain according to Eq.(9) and for whom also Eq.(10) is satisfied.

In case the seller ships the item the buyer's actual transaction gain is given by: $g^{\text{T}}(B) = \text{Val} \otimes \text{Pr}c(S)$. In case the seller does not ship the item the buyer's loss is given by: $g^{\text{T}}(B) = \otimes \text{Pr}c(S)$. In case no deal satisfies Eq.(10) the buyer will not transact in that particular round and the gain is zero.

After every session the sellers and buyers will compute the profit made as a result of trading during that session.

3.2 Market Strategies

Each seller S is initialised with selling price $\text{PrC}(S) = 100$, honesty $\text{Hst}(S) = 0.9$, and base rate reputation parameters $(r_{\text{base}}, s_{\text{base}}) = (9, 1)$. Each buyer B is initialised with risk aversion $\text{Avs}(B) = 0.9$. These parameter values (except the base rate reputation) are subject to change which represents the agents' strategies for optimising their profit.

Seller Mutation. For each seller the session gain $g_k^S(S)$ will be compared with the previous session's gain $g_{k-1}^S(S)$. The seller parameter values that generated the highest gain are used as basis for defining new and possibly better parameters which will be used in the subsequent session $k + 1$. The new parameters are defined as follows:

The price $\text{PrC}(S)$ is increased by Δ_{PrC} with probability $\text{Prob}_{\text{PrC}}^+$, decreased by Δ_{PrC} with probability $\text{Prob}_{\text{PrC}}^-$, and remains unchanged with probability $\text{Prob}_{\text{PrC}}^=$, where:

$$\text{Prob}_{\text{PrC}}^+ = 0.33, \quad \text{Prob}_{\text{PrC}}^- = 0.33, \quad \text{Prob}_{\text{PrC}}^= = 0.34 \quad \text{and} \quad \Delta_{\text{PrC}} = \$1 > \quad (11)$$

The honesty $\text{Hst}(S)$ goes up by Δ_{Hst} with probability $\text{Prob}_{\text{Hst}}^+$, goes down by Δ_{Hst} with probability $\text{Prob}_{\text{Hst}}^-$, and remains unchanged with probability $\text{Prob}_{\text{Hst}}^=$, where:

$$\text{Prob}_{\text{Hst}}^+ = 0.33, \quad \text{Prob}_{\text{Hst}}^- = 0.33, \quad \text{Prob}_{\text{Hst}}^= = 0.34 \quad \text{and} \quad \Delta_{\text{Hst}} = 0.02 > \quad (12)$$

Seller Intelligence. Under certain conditions the seller mutation is overruled by a set of simple rules.

1. If a seller has not conducted any transactions during the last session, the selling price $\text{PrC}(S)$ is decreased by Δ_{PrC} , and the honesty $\text{Hst}(S)$ is increased by Δ_{Hst} .
2. The minimum selling price is zero, thus allowing sellers to transact with loss in order to regain reputation and get back into the market.
3. The maximum and minimum honesty is 1 and 0 respectively.

Buyer Mutation. For each buyer the session gain $g_k^S(B)$ is compared with the previous session's gain $g_{k-1}^S(B)$. The buyer risk aversion that generated the highest gain are used as basis for defining a new and possibly better risk aversion which will be used in the subsequent session $k + 1$. The new risk aversion is defined as follows:

The risk aversion $\text{Avs}(B)$ is increased by Δ_{Avs} with probability $\text{Prob}_{\text{Avs}}^+$, is decreased by Δ_{Avs} with probability $\text{Prob}_{\text{Avs}}^-$, and remains unchanged with probability $\text{Prob}_{\text{Avs}}^=$, where:

$$\text{Prob}_{\text{Avs}}^+ = 0.33, \quad \text{Prob}_{\text{Avs}}^- = 0.33, \quad \text{Prob}_{\text{Avs}}^= = 0.34 \quad \text{and} \quad \Delta_{\text{Avs}} = 0.01 > \quad (13)$$

Buyer Intelligence. Under certain conditions, the buyer mutation is overruled by a set of very simple rules.

1. If the buyer has not conducted any transactions during the last session, then the risk aversion $\text{Avs}(B)$ is decreased by Δ_{Avs} .
2. If the buyer has made a loss in the previous session, then the risk aversion $\text{Avs}(B)$ is increased by Δ_{Avs} *the number of times the item was not received*.
3. The maximum and minimum risk aversion is 1 and 0 respectively.

4 Simulation Results

We ran 3 different simulations in order to compare the performance of market with and without the beta reputation system.

We simulated a market without a reputation system by defining a feedback longevity value $\lambda = 0$ which means that all feedback is immediately forgotten before it contributes to the reputation rating. We simulated a market with a reputation system by setting $\lambda = 0.99$ which causes feedback to be slowly forgotten. With this longevity value, the maximum reputation rating is also 0.99. We also simulated a market with $\lambda = 1$ in order to study the effect on the market when feedback is never forgotten. All 3 simulations ran for 1000 sessions and the graphical illustrations presented below are based on data obtained from those 3 simulation runs.

4.1 Honesty and Reputation

In this section we look at how well the reputation rating follows the honesty of an individual seller. The honesty represents the sellers personal and private state of mind whereas his public reputation rating can be interpreted as an estimate of his honesty. Idealistically the reputation rating should be equal to the honesty, but that can never be achieved. The reliability of the reputation system can be measured by how well the reputation rating follows the real honesty.

Fig.3 illustrates honesty and reputation rating for Seller 1 in a market without feedback.



Fig. 3. Honesty and reputation rating for Seller 1 where $\lambda = 0$

Because all feedback is immediately forgotten the reputation rating remains constant at the base rate 0.9. This allows the sellers to reduce their honesty, thus to be more fraudulent, without suffering a loss of reputation. This causes the market to deteriorate.




Fig. 4. Honesty and reputation rating for Seller 1 where $\lambda = 1$

Fig.4 below illustrates a market where $\lambda = 1$, and where the rating suffers.

In Fig.4 it can be seen that the reputation rating follows the honesty fairly well in the beginning of the trace, but that it deviates more and more further out in the trace. When feedback is never forgotten, the reputation rating becomes less sensitive to new feedback. In the end it becomes possible to reduce honesty without suffering a loss in reputation in the short term. The simulation shows that the seller exploits this effect to reduce honesty and become fraudulent, thus causing the market to deteriorate.

Fig.5 illustrates the evolution of the honesty and reputation rating for seller 1 in a market with diminishing feedback, i.e. with feedback longevity factor $\lambda = 0.99$.

It can be seen that the reputation rating tracks the honesty fairly well. As a result, when honesty for example is reduced it will have an effect on sales in two ways: 1) it will lead to short term increased profit for the seller because some items are not shipped, and 2) it will lead to long term reduced sales because the reputation rating will suffer relatively quickly. The balance between these two effects keeps the honesty within a range of 0.75 - 1.0 during the simulation. This reflects a much more healthy market than when all feedback is either immediately forgotten ($\lambda = 0$) or never forgotten ($\lambda = 1$).

Fig.6 illustrates the evolution of the honesty and reputation rating for one seller with diminishing feedback during the sessions 50-150.

The figure which represents a close zoom-in on Fig.5 is included to show in detail how the reputation rating follows the honesty when the longevity factor $\lambda = 0.99$. It can be seen that during sessions 65-70 when the honesty is 1.0 the reputation approaches its maximum before dropping again.

The fact that the reputation rating follows the honesty in a rather erratic way is of course due to the fact that the reputation rating is a result of feedback which is a result of a probabilistic decision to ship or not to ship in individual transactions. It is thus

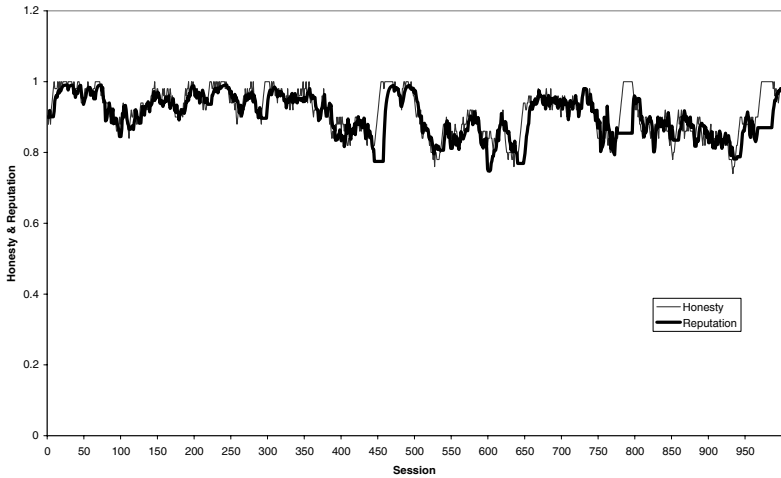


Fig. 5. Honesty and reputation rating for Seller 1 where $\lambda = 0.99$

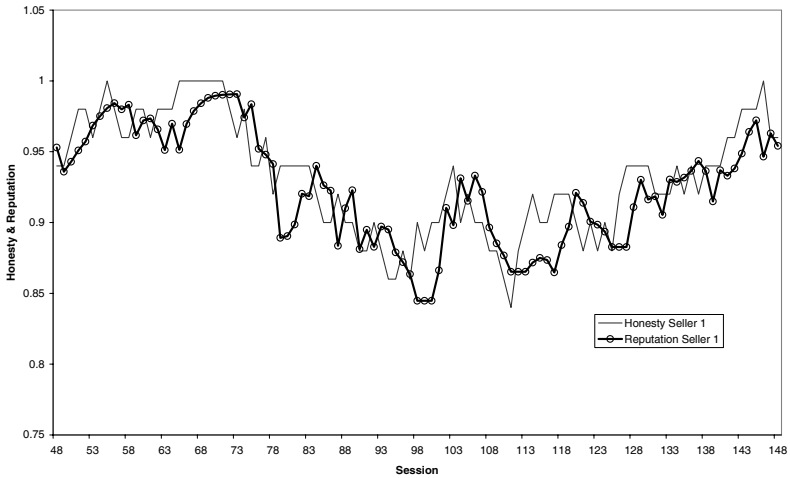


Fig. 6. Honesty and reputation rating for Seller 1 during sessions 50-150 where $\lambda = 0.99$

possible that even though the honesty is 0.90 the seller might for example ship only 5 out of 10 items in any given session.

4.2 Snapshot of Seller Strategy, Rating, and Transactions

The seller strategy is defined by the honesty and the selling price. The tracing of these two parameters together on the same plane can be called a “genetic walk” if we consider the parameters as the “genes” that define the nature of the seller. In Fig. 7 below we have traced the genetic walk for all three sellers during the sessions 600-620.




Fig. 7. Seller strategy evolution during sessions 600-620 where $\lambda = 0.99$

It can be observed that in this period seller 1 and 2 are competing with comparable honesty levels whereas seller 3 operates with a lower honesty level. The resulting reputation rating for the 3 sellers in the same period can be seen in Fig.8 below.




Fig. 8. Seller reputation rating during sessions 600-620 where $\lambda = 0.99$

It can be observed that seller 3 has the lowest reputation rating overall, and that the reputation ratings of sellers 1 and 2 in fact are crossing each other at session 613. The combination of selling price and reputation rating determines the number of transactions

that each seller will get during a session. Fig.9 shows the number of transactions of each seller in the same period.




Fig. 9. Seller transactions during sessions 600-620 where $\lambda = 0.99$

As expected it can be seen that seller 3 has the lowest number of transactions due to his relatively low reputation rating and high selling price.

4.3 Market Honesty

Market Honesty measures the number of shipped transactions divided by the number of total number of transactions. We will measure the market honesty for each session and trace the honesty as a function of session number.

In Fig.10 which illustrates the market honesty when feedback is ignored, it can be seen that the sellers notice that their base rate reputation rating of 0.9 is not influenced by cheating, so that they in fact can improve their profit by reducing the honesty level. The increase in market honesty (around session 550) can be explained by the buyers getting more risk averse and thus buying less, resulting in reduced business for all the sellers. By increasing the honesty the sellers will get more business and increase profit, and the seller mutation process is somehow able to implement this strategy. Even though the general market honesty increases in the period after session 550, it can be seen that there are brief periods of very poor market honesty, which never occurred in the period before session 550. The market around session 800 thus seems to be in a worse state than around session 200, even though the general market honesty level is equal at both instants.

In Fig.11 which illustrates the market honesty when feedback is permanent, i.e. $\square = 1$, it can be seen that the honesty steadily deteriorates. The difference between this and the previous graph (where feedback was immediately forgotten) can be explained




Fig. 10. Market honesty where $\lambda = 0.0$




Fig. 11. Market honesty where $\lambda = 1.0$

by the fact that the sellers are unable to improve business by increasing honesty because of the stiffness of their reputation ratings. Due to their myopic view of the world (only the two last transactions matter) the sellers are unable to keep increasing their honesty long enough to notice that it influences their reputation in a positive way.

Fig.12 above illustrates the market honesty with feedback longevity $\beta = 0.99$. It can be seen that when the feedback is slowly forgotten the market honesty remains significantly higher than is the case when feedback longevity $\beta = 0$ or $\beta = 1$.




Fig. 12. Market honesty where $\lambda = 0.99$

4.4 Profit

Fig.13 illustrated the average profit of sellers and buyers for each session in a market where the feedback longevity $\lambda = 0.99$.




Fig. 13. Profit of sellers and buyers in the market where $\lambda = 0.99$

Two effects can be observed:

- In general the curves are mirrors of each other, meaning that when sellers make more profit over a period, then the buyers make less profit, and vice versa. The periods

of high seller profit can be described as seller's market, and the periods of high buyer profit as buyer's market. One explanation of why the simulation has periods of seller's market and buyer's market is there is an unstable balance between the sellers' and buyers' opposing strategies for maximising profit.

- There are brief periods when there is a sharp profit drop on both seller and buyer side. This is obviously caused by a sharp drop in the number of transactions caused by the sellers' reputation rating dropping below the buyers' risk aversion so that the buyers decide not to transact. These periods are brief because the sellers are programmed to reduce price and increase honesty and the buyers are programmed to reduce risk aversion when no transactions occur, so that the market quickly picks up again. These sharp drops in profit often initiate a period of buyer's market.

5 Discussion and Conclusion

In our study we simulated a market of selling and buying agents where the buyers provide feedback about the sellers after each transaction. We were able to demonstrate some fundamental properties of markets of selling and buying agents:

1. A market without reputation system where strangers transact with each other without knowing each others past behaviour record will be unhealthy.
2. A market with a reputation system in place has the potential to become healthy.
3. A market with a reputation system can only become healthy when the record of past behaviour is slowly forgotten.

Conclusion (3) above intuitively makes sense because most people would judge it unreasonable to put equal weight on behaviour observed in a distant past as on behaviour observed yesterday. It is interesting to notice that the Feedback Forum used on eBay¹ uses a mixture of permanent and forgetting indicators. The overall number of positive minus the overall number of negative feedback is presented together with the number of positive, negative and neutral feedback from the past 6 months, from the past month and from the past 7 days. The reputation rating derived by the beta reputation system represents similar information in a much more condensed form which can be used directly for decision making.

The buying agents in our simulation have a strong ability to forgive sellers' sins of the past, even too strong one might say. For example when a buyer has been cheated, he provides a negative feedback, increases the risk aversion by $A_{vs}(B)$ and a part from that could perfectly well trade with the same dishonest seller in the next transaction as long as the buying criteria are satisfied.

Typical sayings are that it takes many good impressions to cancel out one bad impression, that you never get a second chance to make a first impression, and that once your reputation is lost it is almost impossible to regain. Our trading agents certainly have a significantly more positive view of the world than these proverbs suggest. It thus indicates that the sellers are able to maintain a lower honesty and still survive in the market than would have been possible in the real world.

¹ eBay is the most successful electronic market that uses a reputation system. (<http://ebay.com>)

The buyer side in our simulation consists of only 10 buyers, and even when they have been ripped off by dishonest sellers and their risk aversion stops them from transacting for a few sessions they quickly return to the market as the eternal optimists they are. In the real world, buyers who have been seriously ripped off might leave the market and never return. In the real world there is also a steady stream of newcomers to the market whereas our simulation used a fixed set of players. We deliberately gave the buying agents a positive world view in order to have a functioning market that can produce transaction data. We also thought this was a simple way to implement the combination of pessimistic players who leave and new players who arrive. It would be interesting so observe how the market would be influenced by giving the buyers a more human pessimistic attitude, as well as have a flux of leaving and arriving players.

The simulation results presented here are the first obtained with our market implementation, and we plan to continue our study by changing the market parameters and by analysing other aspects of the market in order to learn more about the ability of reputation systems to influence markets.

The real benefit of reputation systems will emerge as software agents evolve into intelligent entities that are able to automatically make decisions and perform transactions on behalf of their human masters. We believe that software agents will have characteristics of social entities which are able to follow strategies for interactions with other agents for mutual benefit while resisting exploitation by malicious agents. Reputation systems can thus be seen as “navigation instruments” for software agents to find optimal transaction partners and to avoid dangerous encounters. In a nutshell, the goal of reputation systems can be defined as that of optimising the quality of the market and the gain of individuals.

References

1. Gary Bolton, Elena Katok, and Axel Ockenfels. How Effective are Online Reputation Mechanisms? Discussion Papers on Strategic Interaction 25-2002, Max-Planck-Institut, 2002.
2. George Casella and Roger L. Berger. *Statistical Inference*. Duxbury Press, 1990.
3. Consumers International. *Consumers@shopping: An international comparative study of electronic commerce*. Consumers International’s Programme for Developed Economies and Economies in Transition, <http://www.consumersinternational.org/campaigns/electronic/e-comm.html>, September 1999.
4. A. Jøsang and R. Ismail. The Beta Reputation System. In *Proceedings of the 15th Bled Electronic Commerce Conference*, Bled, Slovenia, June 2002.
5. US FTC. Boom in E-Commerce Has Created Fertile Ground for Fraud. US Federal Trade Commission, <http://www.ftc.gov/opa/2001/05/iftestimony.htm>, May 2001.