Department of Computer Science
Series of Publications C
Report C-2005-10

Bayesian game theory in practice: A framework for online reputation systems

Petteri Nurmi

University of Helsinki
Finland

Bayesian game theory in practice: A framework for online reputation systems

Petteri Nurmi
Department of Computer Science
P.O. Box 68, FI-00014, University of Helsinki, Finland
petteri.nurmi@cs.helsinki.fi

Department of Computer Science, University of Helsinki
Technical report, Series of Publications C, Report C-2005-10
Helsinki, March 2005, 22 pages

Abstract

As the number of online auction sites has increased the interest towards providing reliable summaries, reputations, about the past behaviour of the sellers has risen. However, many of the existing approaches are based on various heuristics and properly evaluating the systems is difficult.

To improve the situation, we first introduce a game-theoretical model for reputation in online auctions. Furthermore, we show how mixture models can be used as a proper framework for integrating information from various sources into the model. We also discuss practical aspects and evaluate an experimental online reputation system that has been implemented based on the proposed framework. Our results show that, using this framework, the number of successful transactions increases significantly (in statistical sense) and that the framework is robust for changes in the behaviour of the participants.

Computing Reviews (1998) Categories and Subject Descriptors:

G.3 Mathematics of Computing - Probability and Statistics - Statistical computing

I.2.6 Computing Methodologies - Artificial Intelligence - Learning

J.4 Computer applications - Social and Behavioural Sciences - Economics

General Terms:
Algorithms, Economics, Experimentation, Theory

Additional Key Words and Phrases:
Reputation systems, game theory, mixture models

1 Introduction

Many online portals, e.g. eBay [10], Amazon [3] and Epinions [13], provide their users the opportunity to arrange auctions in a large online community. Unfortunately this kind of system may easily attract selfish users who attempt to utilize the system for their personal benefit. To cope with the situation, the customers are given summaries about the past behaviour of the sellers on which to base their buying decisions. If the mechanism giving the summaries is designed properly, the summaries are such that they discourage selfish behaviour and reward cooperation. In game theoretic terms the summaries, i.e. reputations, are *incentives* [17] as they encourage the users of the system to cooperate, i.e. play by the rules. Reputation can be seen as a punishment-based incentive scheme as the expected utility of a seller decreases if (s)he misbehaves. On the other hand, empirical studies on eBay have shown that sellers with high reputation often are offered higher prices [24] and thus reputation is also a reward-based incentive.

To give practical insights to the problem, we consider the eBay auction site, which can be seen as a economic market. In the eBay auction market participants are divided into (possibly overlapping) sets of customers (buyers) and sellers. The customers can buy items using the following scheme: after the price has been agreed upon, the buyer transfers the agreed sum to the account of the seller who then delivers the items that were part of the agreed transaction. However, if the buyers have no additional information about the seller, the optimal strategy of the seller is to always cheat, i.e. to keep the money and not to send anything to the buyer. Another possibility to cheat is to send an item of lower quality than the originally agreed upon. The reputation scheme adopted by eBay attempts to overcome cheating by allowing the buyers to give positive (+1), neutral (0) or negative (-1) feedback about the seller after a transaction has been concluded and showing the cumulative sum of the feedback values to all buyers. The buyers can also post short comments about the sellers and these comments are shown to all customers. However, empirical analysis by, e.g., Resnick et al. has shown that even though the market works relatively well, it still has many shortcomings [41]. The main reason for these shortcomings is that theoretical optimality of the eBay reputation scheme is based on unrealistic assumptions about the behaviour of the customers [9].

The concept of reputation has various applications in diverse fields such as economy, scientometry, computer science, sociology and evolutionary biology [36]. Thus also the problem of finding a proper theoretical model has already been addressed from various viewpoints. Still, none of the existing models seems to sufficiently explain the behaviour that researchers have observed in, e.g., eBay. We attempt to improve the situation by proposing a new game theoretic model for eBay-like online auctions. In the model, the players are not seen as selfish rational beings, but ethically rational players that consider the norms of the society in their decision making. This interpretation supports empirical observations [40] and is inline with the current view of human decision making, according to which the rationality of humans is bounded, i.e. optimal up to limited computational capabilities, [19]. Our model is further discussed in Sections 4 and 5.

In the proposed model we maintain a single estimate of the reliability of a seller. However, as information often originates from multiple sources, we need a way to integrate different information sources in a flexible and theoretically justifiable way. The solution is to use *mixture models* [32] which additionally make it possible to calculate personalized reputation values for different buyers. This topic is further discussed in Section 6.

Although the model that we propose is theoretically optimal, it is still somewhat insufficient in practice for reasons that are discussed later. To boost the performance of implementations based on the proposed framework, we introduce in Section 7 some heuristics for providing additional information for the buyers.

To further support our arguments, we implemented an online reputation system based on the proposed framework. In addition, we implemented a similar system that uses the eBay reputation mechanism and compared the performance of the two systems. The results of the experiments are discussed in Section 8. Finally, in Section 9 we conclude the paper and discuss future work. However, before going into details about the model, we introduce some background information and discuss related work.

2 Background

Bayesian probability theory

Bayesian probability theory is a statistical theory of making statements about uncertain events θ. Initially events of interest are assigned a prior belief $p(\theta)$ which reflects existing knowledge about the event and the problem area. Later, as new information $\mathcal{D}$ becomes available, the subjective beliefs are updated using the Bayes' rule

$$\mathit{posterior} = p(\theta|\mathcal{D}) = \frac{p(\mathcal{D}|\theta)p(\theta)}{p(\mathcal{D})} = \frac{\mathit{likelihood} \times \mathit{prior}}{\mathit{normalizer}}. \tag{1}$$

The likelihood term $p(\mathcal{D}|\theta)$ measures the probability of seeing particular realizations of the event θ whereas the normalizer $p(\mathcal{D})$ is used to ensure that the values of $p(\theta|\mathcal{D})$ sum up to one and thus define a proper probability distribution. After the update, the values of the posterior $p(\theta|\mathcal{D})$ are used as the new priors $p(\theta)$.

Consider the case where the likelihood $p(\mathcal{D}|\theta)$ is given by a binomial distribution, e.g. the observations are a sequence of independent coin tosses. If we assign the prior probabilities using a *Beta distribution*, it can be proven that the resulting posterior beliefs are also Beta-distributed (see e.g. [18]). In this case we say that the Beta distribution is the conjugate prior of the binomial distribution and thus we can model the beliefs about events that follow a binomial distribution using a Beta distribution. In addition, this results in an efficient update rule for the posterior beliefs.

Bayesian games

Bayesian games [17, 21, 22, 23] are a combination of game theory and probability theory that allow taking incomplete information into account. In Bayesian games each player is allowed to have some private information that affects the overall game play, but which is not known by the others. However, others are assumed to have beliefs about the private information. These beliefs are represented using probability distributions and they are updated using the Bayes' rule (see Equation 1) whenever new information is available. Optimality in this setting requires that the players act optimally according to their beliefs and their private information.

3 Related work

Modelling

The successfull approaches to modelling reputation are based on game theory and use either Bayesian games or *evolutionary game theory* [4, 48]. An exhaustive overview of different game theoretic models for reputation can be found at [17].

Initially the concept of reputation was discussed by Akerlof when he introduced the *market of lemons* problem [1]. In the problem a person is assumed to buy a car from the market and the car can be either a good car or a lemon (a car of bad quality). The buyer must pay for the car in advance, but (s)he will know the quality of the car only afterwards. According to game theoretic arguments, the rational seller should sell a lemon because this will maximize the seller's personal gain. Similarly, if this model is replayed a discrete number of times, the optimal strategy is selling a lemon at each instance of the game (stage). However, in reality the behaviour of this kind of markets is usually different and thus we need to adopt different optimality conditions for the players.

In modelling reputation the most important optimality condition is called the *sequential equilibrium* and it states that the actions of the buyers should be optimal given the previous game play and the beliefs the buyers have about the behaviour of the seller [27]. Moreover, sequential equilibrium assumes that the beliefs of the buyers are sequentially rational which means that they should eventually converge to represent the correct probability distribution of the types (sell a good car or a lemon). The first proper game theoretic models, introduced by Kreps and Wilson [26] and by Milgrom and Roberts [34], were based on sequential rationality in repeated Bayesian-games. In repeated Bayesian-games a single realization, stage, of the game is modelled using a Bayesian game and this stage game is repeated over time. Moreover, the beliefs of the players are updated after each realization. The problem with this approach is that Bayesian games only allow simultaneous moves, i.e. both the sender and buyer must decide their actions at the same time. The generalization to games with non-simultaneous moves was introduced a few years later by Fudenberg and Levine [16] who allow the stages to be *extensive games* [28]. This work also forms the basis of our work.

Also various modifications have been introduced. For example, Alt et al. consider a model where leadership is incorporated into the model so that leaders with weak resources can obtain a haegemonic position if their reputation is good enough [2]. Other work includes, e.g., the work by Ely and Välimäki [12], who show that reputation can sometimes be unadvantageous for a player. An analysis of when reputation is bad is presented by Ely et al. [11].

In the context of evolutionary game theory there also has been much work regarding reputation. Most of the evolutionary approaches are based on the *iterated prisoner's dilemma* (e.g., [17]). The underlying theory is that the players are motivated to cooperate if their non-cooperative actions result in a loss of gain in the future. This effect was introduced by Axelrod, who calls it the *shadow of the future* [4]. There are two variations of the basic model that differ in whether the prisoners' dilemma is repeated an infinite or a finite number of times. Both variants have also been studied with computer simulations, the infinite case in, e.g., [25] and the finite case in, e.g., [30].

Practical solutions

The properties of eBay have been widely studied within computer science and economics. There has been both experimental evaluations and theoretic approaches that aim to identify weaknesses in eBay's reputation scheme or that attempt to model some peculiarities of the observed market behaviour. Of special interest is the empirical work conducted on eBay by Resnick et al. [41]. According to them, the marketplace works rather well even though the mechanism itself has many shortcomings such as easy access to new pseudonyms [15], lack of feedback and guaranteeing the trustworthiness of provided feedback [35]. Other empirical studies on eBay have shown, e.g., that the prices the sellers are paid usually increase as reputation grows [24].

In other online communities the work has mainly concentrated on proposing new mechanisms. For example the REGRET system [42] uses the Beta-distribution for maintaining reputation values (see Section 2). Reputation systems have also been developed for the users of rating systems. One such scheme is presented in [7].

Another important practical aspect is how to aggregate information from other agents in a reliable way into the used reputation model. A major research area in this field is using social networks as a source of information. This has been studied in, e.g., [39, 43].

Other uses of reputation have emerged in different fields of computer science. Especially, specialized communication networks, such as P2P networks [8, 14, 20, 31, 44, 47] and wireless ad hoc networks [5, 6, 33], have used the concept of reputation to optimize the behaviour of the participants. Due to the distributed nature of the solutions, the approaches in these fields differ from the web market area slightly as reliability of information must be taken into account before different sources of information are integrated.

4 The reputation game

In this section we define the structure of the reputation game that the buyers and sellers are playing in an online auction site. Our game model differs from previous work by incorporating *reciprocal altruism* [46] into the model. According to reciprocal altruism, part of the sellers' cooperation (delivering the items) results from the fear of being cheated later on. The overall effect this fear has on the actions of a seller varies from one seller to another, but the buyers can try to estimate the overall effect and to base their decisions on these estimates. However, before going into details about the model, we need to define some notation.

Let $\mathcal{M}$ be an online marketplace with a finite set of buyers N and a finite set of sellers M. An arbitrary buyer is indexed using the variable i whereas sellers are indexed with the variable j. The online nature of the market determines the transaction policy of the system which is assumed to be similar as in eBay. Thus after the seller and buyer have reached an agreement about the price (and items), the buyer transfers the agreed amount of money to the seller's bank account and then the seller delivers the agreed items to the buyer. The transferred money amount is called the cost of the transaction and the variable c_{ij} is used to denote it.

We assume that each seller j has a discrete representation for time so that time step t_j^k denotes the k^{th} time j is selling item(s) to the market. The cost of the k^{th} transaction is denoted by $c_j(t_j^k)$. At each time step the seller plays a simple game against some buyer i. The games played at individual time steps are called *stage games* and we assume that their structure is as in Figure 1. The structure is discussed more thoroughly later in this section. The collection of stage games is a *supergame* [45] and we call this collection the *reputation game* between a set of buyers $N\backslash\{j\}$ and a seller j.

The stage games are assumed to be Bayesian games with asymmetric information, i.e. one of the players has more knowledge about the game than the others. In the online auction model the asymmetry results from the market rules as the buyer does not know whether the seller is *HONEST* (delivers the items) or *DISHONEST* (cheats). Thus the seller can have two modes of behaviour which are called *types* [17] in game theoretic terms. The *HONEST* seller never cheats whereas the *DISHONEST* seller can cheat by not submitting anything or by submitting an item of lower quality. Furthermore, we assume that each seller acts as a *HONEST* merchant in a fraction of the transactions and that this fraction is determined by the seller. We use the variable θ_j to denote the fraction of transactions, where the seller j acts honestly and thus the probability distribution governing the behaviour of the seller is determined by $p(HONEST) = \theta_j$ and $p(DISHONEST) = 1 - \theta_j$.

If we ignore the effect of the shadow of the future (see Section 3), the best value for parameter θ is $\theta_j = 0$. However, if we take the shadow of the future into account, the selection of the parameter becomes more complicated. The process according to which this parameter is actually chosen is not known, but for our purposes it is not even necessary to know the details of this process[1]. However, as higher values of the parameter θ_j result in higher loss-rates in individual stage games, we call the parameter the *reciprocity rating* of the seller.

Initially, i.e. at time step t_j^1, the buyers have no information about the reciprocity rating of seller j. However, as more and more transactions are concluded, the amount of available information increases and this should affect the buying decisions of the customers. We assume sequential rationality and initially assign a belief distribution μ over the possible parameter values and update these beliefs using the Bayes' rule as new information is obtained. According to Bayesian decision theory the best estimate of the seller's honesty at an arbitrary time step t_j^k is given by the mode of the belief distribution which is defined as

$$\hat{\theta}(t_j^k) = \arg\max_{\theta} \mu(\theta). \qquad (2)$$

Assuming (for now) that the underlying true distribution of the types remains fixed, the beliefs of the buyers converge to the true distribution. If we can furthermore prove that the individual stage games have optimal strategies, using similar arguments as in [38], it follows that the model admits a *perfect Bayesian equilibrium.* We next show that the stage games indeed have optimal strategies by discussing their structure more thoroughly. For this we consider

[1]For further details see the discussion about inductive theory and deterrence theory in [45].

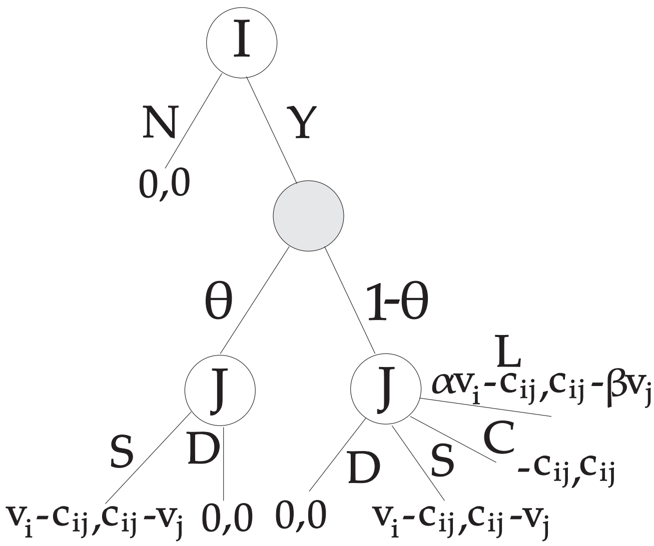

Figure 1: Single stage of the reputation game.

the k^{th} transaction taking place between a seller j and some buyer i. For notational simplicity, we drop the time indices (t_j^k) in the rest of this section.

We assume that both the buyer and the seller have a valuation for the item(s) that are part of the transaction. The valuations are denoted by v_i and v_j and they represent how much the item(s) are worth to the participants. We assume that the participants have already agreed on the price c_{ij}, but that the money has not been transferred. In this case we are in the root node of the game-tree depicted in Figure 1. According to the market rules, the first one to perform an action is the buyer (capital I in the figure). The buyer can still cancel the transaction by not transferring the money (N) or (s)he can go ahead with the transaction (Y). If the buyer hesitates and cancels the transactions, the resulting situation is the same as in the beginning of the stage and thus the utilities of the participants are zero.

The more interesting scenario occurs when the buyer decides to go ahead with the transaction. If the seller delivers the correct items (S), the buyer loses c_{ij} units of money, but receives the item(s) which have value v_i for the buyer. Thus the utility of the buyer is $v_i - c_{ij}$. Similarly, the seller gains c_{ij} units of money, but loses the item(s) which had the seller the value v_j. Thus the resulting utility for the seller is $c_{ij} - v_j$.

On the other hand, the seller's valuation of the item(s) may change over time so that when (s)he gets the money the difference $c_{ij} - v_j$ is negative and the seller can decline (D) and return the money in which case the utilities of the participants are again zero.

The analysis thus far concludes the behaviour of the honest seller. However, if the seller is dishonest, (s)he has the additional options of not submitting any item(s) (cheating, C) and submitting item(s) of lower quality (L). If the seller decides not to submit anything, the buyer has lost c_{ij} units of money without getting anything. Thus her(his) utility is $-c_{ij}$. Similarly, the seller has obtained c_{ij} units of money without losing anything and thus the seller's utility is c_{ij}.

The more difficult case occurs when the seller decides to submit item(s) of lower quality (L). In this case we can assume that both the buyer and seller have a subjective estimate of the level of quality of the item(s). We use the variable α ($0 \leq \alpha \leq 1$) to denote the buyer's estimate and the variable β ($0 \leq \beta \leq 1$) to denote the seller's estimate. Thus the products of lower quality have a discounted value of αv_i for the buyer. In this case the buyer loses c_{ij} units of money and gains an item with value αv_i which results in an overall utility of $\alpha v_i - c_{ij}$. Similarly, the lower quality item has a value βv_j for the seller and, analogously, the utility of the seller is given by $c_{ij} - \beta v_j$.

Clearly in a single stage game the optimal strategy of the honest seller is to go ahead with the transaction (S) if $c_{ij} - v_j \geq 0$. Similarly, the optimal strategy of the dishonest seller is to cheat whenever $c_{ij} \geq 0$.

The optimal strategy of the buyer is a bit more complicated. If the buyer knows the underlying probability distribution of θ_j, (s)he could act optimally by calculating the expectation of the dominant strategies. However, as the true distribution is unknown, the only possibility for the buyer is to act optimally according to her(his) beliefs. Thus the optimal strategy of the buyer is to go ahead with the transaction if and only if $\hat{\theta}(v_i - c_{ij}) + (1-\hat{\theta})(-c_{ij}) = \hat{\theta} v_i - c_{ij} \geq 0$. As the values of the variables v_i, v_j and c_{ij} are determined in the beginning of the game, the stage game has equilibrium strategies which are determined according to the discussed rules.

5 Sequential rationality in the reputation game

The theoretical model discussed in the previous section assumes that we maintain an estimate of the honesty of the seller, but gives no practical insights on how to do this. The solution is to look at the problem from a statistical viewpoint and to consider the sequence of transactions as if it was a sample from a binomial distribution, i.e. a sequence of independent coin tosses, with a bias parameter θ. Thus we can estimate the reciprocity rating using Bayesian parameter estimation of the binomial distribution. Furthermore, assuming that we use a smoothed Beta distribution as the prior distribution, the formula for calculating the posterior parameter estimate actually coincides with the maximum likelihood estimate.

Instead of assuming each transaction as a binary outcome, we take into account the cost of the transaction. Thus instead of estimating the proportion of successful transactions, we estimate the proportion of the cost involved in successful transactions. This way the punishments and rewards of the system make exploiting the system more difficult[2].

Let $\mathcal{A}_j$ be the set of all transactions where the seller has been j. Of these a subset $\mathcal{C}_j$ has been successful, i.e. the seller was honest. We define the variable $\hat{\theta}_j$ to be the ratio of successful transactions and all transactions, i.e.

$$\hat{\theta}_j = \frac{\sum_{\mathcal{C}_j} c_{ij}}{\sum_{\mathcal{A}_j} c_{ij}}. \tag{3}$$

The Equation 3 is the maximum likelihood estimate of the parameter θ and it satisfies sequential rationality.

If we use Bayesian parameter estimation and assign a Beta-prior on θ, the resulting posterior estimates are also Beta-distributed, as discussed in Section 3. The Beta-distribution is a two parameter distribution whose parameters are denoted by a and b. The parameter a measures the number of successes and, analogously, b measures the number of failures. The bias estimate of the underlying binomial distribution is given by the mode of the Beta distribution which is given by

$$\frac{a-1}{a+b-2}. \tag{4}$$

Assuming standard prior smoothing is used, we set initially $a = b = 1$. If the k^{th} transaction was successful, we update the parameter a by setting $a^{k+1} = a^k + c_{ij}$ and if the seller cheated we update the parameter b by setting $b^{k+1} = b^k + c_{ij}$. Substituting the resulting values into Equation 4 results in the following Equation for the mode of the posterior distribution

$$\hat{\theta} = \frac{(\sum_{\mathcal{C}} c_{ij} + 1) - 1}{(\sum_{\mathcal{C}} c_{ij} + 1) + ((\sum_{\mathcal{A}} c_{ij} - \sum_{\mathcal{C}} c_{ij}) + 1) - 2} = \frac{\sum_{\mathcal{C}} c_{ij}}{\sum_{\mathcal{A}} c_{ij}}, \tag{5}$$

which is exactly the maximum likelihood estimate.

6 Integrating Reputations using Mixture Models

In online reputation systems the feedback about the behaviour of a seller is derived from interactions of various customers. However, as the reliability of a feedback is hard to estimate, individual buyers base their decisions mainly on personal experience and consider the information from others as a confirmation of their own beliefs. Current reputation systems fail to properly aggregate

[2]Currently we ignore the possibility that the buyer intentionally gives negative feedback.

the different information sources so that a personalized reputation rating can be calculated for the individual buyers. Our proposal to this task is to use *mixture models* [32] in which the overall beliefs are represented as a sum of weighted "component" beliefs. The mixture components represent the different sources of information and the weights represent the importance we assign to a particular information source. In mathematical terms a mixture density is defined as

$$p(\theta_j) = \sum_i \frac{\pi_i}{\sum_i \pi_i} p_i(\theta_j), \tag{6}$$

where π_i is the weight of the i^{th} mixture component and p_i is the underlying probability density of that component.

Each of the component densities p_i maintains an estimate of the honesty of a seller. From a theoretical point of view the internal structure of the component densities is not important and they can be, e.g. individual binomial distributions, *hierarchical models* [18] or *Bayesian networks* [37]. In addition, the number of components can be quite freely selected. However, for our purposes it suffices to consider a model that consists of two mixture components. We use p_1 and p_2 to denote the component densities and π_1 and π_2 to denote the mixture weights of the components. According to this model, the overall estimate of the honesty of the seller is given by

$$p(\theta_j) = \frac{\pi_1}{\pi_1 + \pi_2} p_1(\theta_j) + \frac{\pi_2}{\pi_1 + \pi_2} p_2(\theta_j). \tag{7}$$

The interpretation that we propose for the mixture weights is that each weight term represents our estimate of the reliability of information from a particular source and that the values lie in the interval $[0, 1]$. For example, if we trust our own experiences almost completely, we could assign a mixture weight of $\pi_1 = 0.99$ for our personal experiences. Similarly, if we assume that information from others is not very reliable, we could assign a weight term $\pi_2 = 0.20$ for the information provided by others. This gives us the following estimate for the honesty of a seller

$$p(\theta_j) = \frac{0.99}{0.99 + 0.20} p_1(\theta_j) + \frac{0.20}{0.99 + 0.20} p_2(\theta_j). \tag{8}$$

Furthermore, assume that both components are binomial distributions so that the posterior beliefs of the components are Beta distributions with parameters (a_1, b_1) and (a_2, b_2). To give more intuition into our model, we assume that $p_1(\theta) = 0.9$, $p_2(\theta) = 0.18$ and that the mixture weights are as before. In this case the overall honesty estimate is given by

$$p(\theta_j) = \frac{0.99}{0.99 + 0.20} 0.90 + \frac{0.20}{0.99 + 0.20} 0.18 = 0.7790. \tag{9}$$

From our example we can make an interesting observation. Namely, the good personal experiences of the user cancel out bad experiences by other users. This can be interpreted so that although the seller in general cheats, (s)he does not cheat when selling products to a particular user. Moreover, if our transactional history is large enough, this kind of behaviour can provide evidence about social relationships.

As the number of transactions increases, the values of the components concentrate more and more around the true parameter value as the variance of the individual components shrinks. Moreover, if the resulting mixture density is unimodal, i.e. has only one peak, the behaviour of the seller is similar for all users.

The model can naturally be extended to include more mixture components. However, when we consider online markets such as eBay, two components seem to be sufficient. In this case the components correspond to the information derived from the personal experiences and to the information derived from the experiences of the others. Another possible source of information is the relationships in the social network of the user. In the next section we also consider incorporating an estimate of short term behaviour into the model.

Finally, as the probabilities themselves are not very intuitive for the user, we need to map them into more meaningful values. This can be accomplished by applying on the probabilities an arbitrary (weakly) monotonic function $f(\cdot)$, which is continuous on the interval $[0, 1]$. Thus the final reputation values are determined by the value $f(p(\theta))$. Furthermore, if different portals use a similar scheme, but with different mappings f, the reputation values are comparable across different web-portals.

7 Extensions

From a theoretical perspective, the generic framework that we have discussed is sufficient for modelling the problem. However, there are some remaining aspects of the model which are ad hoc by nature and need to be solved using heuristics. First of all, we need to look at how to set the mixture weights and how to adapt them as new information arrives. Secondly, sequential rationality in the model is based on the assumption that the underlying dynamics are stationary which does not hold in practice. Thus we need to discuss the effect of dynamics and how to give estimates about the dynamics of the behaviour. Finally, we shortly discuss how the model can be augmented using Bayesian networks as the mixture components.

Setting mixture weights

The standard way of using mixture models is to gather some data and then use the EM algorithm to learn the model parameter from that data. However, in our task the data arrives in an online manner and thus this approach is not suitable. In addition, as the mixture weights represents confidence levels of information sources, the weight values should increase as new data becomes available.

Our solution model uses exponential growth to achieve the adaptation of mixture weights. We initialize the values of π_i to constant values λ_i. We require that the constants sum up to one, i.e. $\sum_i \lambda_i = 1$. These constants are chosen so that magnitudes of the values reflect prior knowledge about the reliability of the different sources.

After we have initialized the weights, we derive an update term α_i for the weight π_i using the equation

$$\lambda_i \alpha_i^{t_i} = 1.0, \tag{10}$$

where t_i is a positive integer. The updates are then performed by multiplying the current value of the mixture weight π_i with the update term α_i each time new information becomes available.

Intuitively the value of t_i should be smaller for the component representing personal information than for the component representing information from others. In principle these values could be set by a domain expert, but in practice it suffices to use common sense, i.e. to set a small (a magnitude of tens) value for t_1 and much bigger value (a magnitude of thousands) to t_2. The experimental setting uses $t_1 = 10$ and $t_2 = 1000$ and thus personal experience is weighted much more than the information gathered by others.

Optimizing performance under non-stationary dynamics

A drawback of the theoretical model is that if the behaviour of a seller abruptly changes, the convergence of the reputations to a value that correctly represents the new behaviour can be slow. The convergence speed depends on various factors, such as the size of transaction history and the cost of the transactions. We can overcome some of these problems by giving the users as additional information a short term estimate of the behaviour of the seller.

For deriving the short term estimator we weight the costs $c(t)$ of the most recent transactions using the function
$\exp(-n/\mu)$, where the variable $n = 0, 1, 2, \cdots$ represents the n^{th} most recent transaction and μ is a fixed parameter. An example of the shape of this function is given in Figure 2 where we have chosen $\mu = 3.0$. Otherwise the estimator is derived as the mode of a Beta distribution where the parameters

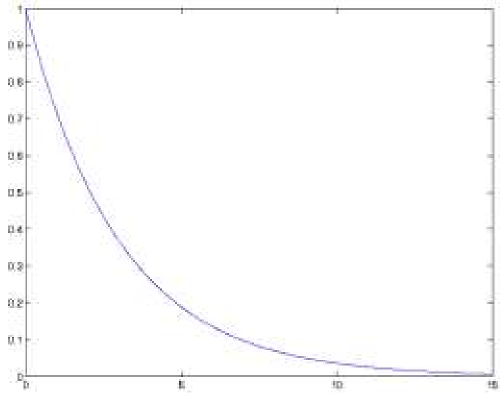

Figure 2: The shape of the function $\exp(-x/\mu)$ with parameter choice $\mu = 3$.

have been estimated from a different subset of informations. Formally the parameter estimate in this case is given by

$$\tilde{\theta}_j = \frac{\sum_{n,\mathcal{C}} \exp(-n/\mu)c(t_{k-n}) - 1}{\sum_{n,\mathcal{A}} \exp(-n/\mu)c(t_{k-n}) - 2}. \tag{11}$$

Bayesian networks for reputation systems

As a final extension we consider the case where the mixture component itself is a Bayesian network. Bayesian networks are directed acyclic graphs that model relationships between variables using probability theory. In the causal interpretation, two (or more) variables are connected through an edge only if there is a direct causal relationship between the variables. According to this interpretation, we can attempt to build a more generic model which attempts to model the relationships different factors have on the overall probability distribution. However, although the Bayesian network approach is more generic, it also suffers from drawbacks. Namely, as the number of nodes in the network increases, the representation of the network requires more parameters and the time required to do the probability calculations increases. The standard way to perform the calculations is to use a message passing algorithm that divides the network into separate partitions and propagates messages among the partitions [37]. Due to space limitations, these algorithms are not further discussed in the paper.

Basically, Bayesian networks can be used to improve reputation systems in two ways. Firstly, by gathering log data from the behaviour of sellers and buyers it is possible to learn a Bayesian network from the data. This network model can then be incorporated into the system and used to "explain" the behaviour of the users. Secondly, we can explicitly add other variables that reflect our prior knowledge about the problem area. For example, when applying this framework for ad hoc or P2P networks, augmented components are required for the integration of information from others due to the distributed nature of the solutions. For example, it is necessary to add an explanatory variable describing the reliability of the source of information. Further discussion about these issues is omitted in this paper.

8 Experiments

To provide empirical evidence about the performance of the Bayesian mixture model approach, we implemented a web-based test application called $B - HONEST$. The application is accessible through the web-site

`http://db.cs.helsinki.fi/u/ptnurmi/`
`cgi-bin/b-honest/index.php`.

In addition, we implemented a similar site that uses the eBay reputation scheme. This version is accessible through

`http://db.cs.helsinki.fi/u/ptnurmi/`
`cgi-bin/ebay/index.php`.

We begin by discussing details about the implemented systems after which we present the experimental setting and the results of our experiments.

Experimental system

Both reputation mechanisms were tested using a similar system. In the implemented system anonymous buyers can buy products from anonymous sellers using virtual money. The sellers are added to the system by the administrator whereas the buyers are real humans. The market rules of the system are the same as has been described earlier, i.e. the buyer first submits the money after which the decision of the seller is made. Currently the only supported actions for the seller are cheating and delivering the (virtual) items.

The decision making process of the seller is, by default, based on stochastic rules which are governed by a parameter θ. The parameter θ serves as the baseline cooperation value, i.e. the smaller the value the more often the seller cheats. The values of θ are otherwise fixed, but the administrator can change the values at any time. The stochastic rules can be overruled using defect and coordinate lists which are determined by the administrator. These lists tell to whom the sellers should always sell and whom to always cheat. In normal reputation systems the identities of the buyers and sellers are unknown to each other, but we added this functionality to simulate relationships in social networks.

Experimental setting

Our test setting consisted of two tests. In both tests a fixed number (6) of persons were required to perform a set of transactions in a simulated market. Three of the users used the $B-HONEST$ system and the three others used the eBay system. The goal of the test subjects was to minimize the sum of costs in unsuccessful transactions.

In the first test, the buyers were required to perform 50 transactions in the market. This setting was repeated three times so that the resulting data set consisted of 900 measurements of which 450 were from $B-HONEST$ and 450 were from eBay. From each transaction we recorded the selected seller, the price of the transaction and the actual outcome (success vs. failure). During a single test run the marketplace had ten sellers with different parameter values θ_j. These values were initialized at the beginning of each test run to random values after which they were kept fixed during the test run. Both systems, of course, had the same parameter values.

In the second test, we gave the reputation values of the buyers' to the test persons and changed randomly the parameters θ governing the behaviour of the sellers. We applied this test setting two times for two sets of sellers. In the first set, we decreased the θ values for seven users and increased them for three users. In the second set we reversed these quantities and decreased the values for three sellers and increased for seven sellers. The goal of this test was to see the change in the buyers' performance as the behaviour of the sellers changes.

As a baseline we used the eBay reputation scheme because, although it is not optimal, it works reasonably. In the first test we used two evaluation metrics: the number of successful transactions and the proportion of sum of successful transactions. For the second test we used only the number of successful transactions.

Results

Although our test setting and the number of tests arranged were not exhaustive by any means, the test results indicate that our reputation scheme works significantly better than eBay's. We begin by discussing the results of the first test.

First of all, we measured the ratio of successful transactions in the 450 first transactions (test one). The results are illustrated in Figure 3. Clearly, our system converges much faster and achieves better overall performance. To further support our claim, we performed standard significance testing and

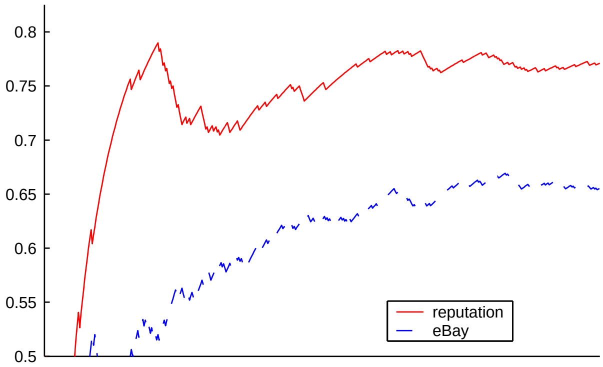

Figure 3: Comparison of the ratio of successful transactions when our model is used (solid) and when eBay system is used (dashed).

calculated the $z - score$ for the difference of the mean performance of the two systems. The test-statistics are given in Table 1 and according to the resulting P-value, the performance increase is statistically significant. Overall, the end performance of the systems differed by 11 percentage units in favour of *B-HONEST*.

SYSTEM	**Variance**	**Mean**
eBay	31.6667	30.6667
B-HONEST	34.8095	36.6667
z-score	-2.8501	
p-value	0.0022	

Table 1: The statistics and test-scores of the experiments.

In real systems, the users are willing to use higher amounts of money only if they trust the sellers and the system. To measure this, we used the formula in Equation 3, i.e. the proportion of costs in successful transactions as the second evaluation metric. Figure 4 illustrates the results of the first test, when this evaluation metric was used. The end perfomance of *B-HONEST* was 78% whereas for eBay it was below 67%. Thus again, our system was around 11 percentage units better. However, although this test did not result in higher differences, it showed that in more realistic situations the performance of our system is very good.

For evaluating the second test we used only the ratio of successful transactions. The results of this are shown in Figure 5. In this case, the difference in performance varied a lot and the minimum difference was around six percentage units whereas the end performance differed actually by 14 percentage units. The main reason for the "poor" initial performance of our system in this test was that users were not completely familiar with our system. Thus as they became more familiar, the system clearly outperformed eBay.

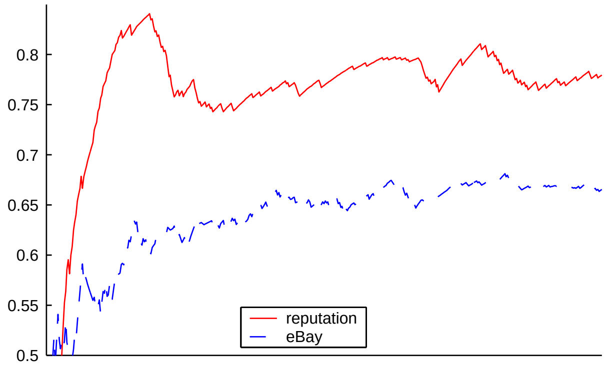

Figure 4: Comparison of the ratio of money involved in successful transactions when our model is used (solid) and when eBay system is used (dashed).

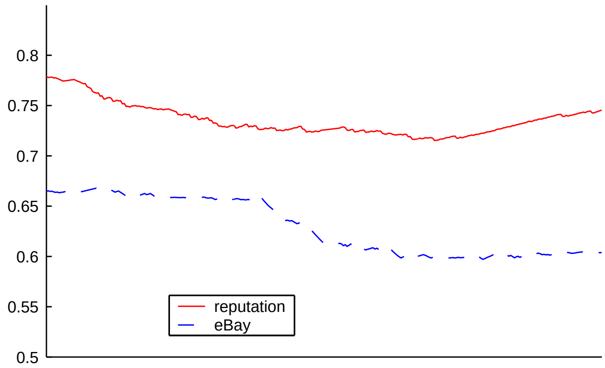

Figure 5: Comparison of the ratio of successful transactions after the behaviour of the sellers' abruptly changed. The solid line represents our model and the dashed line represents eBay.

9 Conclusions and Future Work

In this paper we have presented a theoretical model for modelling reputation in online auctions. In addition, we presented a framework for integrating information from multiple sources. This framework uses mixture models and thus has proper theoretical foundations. Finally, we evaluated our system against the binary reputation scheme used on eBay and, as the results indicate, our system usually outperforms eBay by more than 10% percent.

The work conducted in this paper offers directions for both theoretical and practical future work. Currently we ignored such options as the buyers giving intentionally bad feedback and the effect of discount factors on the selection of the stragegy of submitting an item of lower quality. In practical settings our first goal is to apply the proposed framework to ad hoc and sensor networks. In the field of online auctions, our goal in the future is to consider privacy and security related aspects of our model.

10 Acknowledgments

This work was supported in part by the IST Programme of the European Community, under the PASCAL Network of Excellence, IST-2002-506778. This publication only reflects the authors' views. The author wishes to thank Greger Lindén, Patrik Floréen and Michael Przybilski for reading earlier drafts of the paper and giving valuable comments on the contents. In addition, the author wishes to thank the test persons for performing the experiments.

References

[1] George A. Akerlof. The market for lemons: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, Aug 1970.

[2] James E. Alt, Randall L. Calvert, and Brian D. Humes. Reputation and hegemonic stability. *The American Political Science Review*, 82(2):445 – 466, June 1988.

[3] Amazon. http://www.amazon.com/, 2004.

[4] Robert Axelrod. *Evolution of Cooperation*. Basic Books, Cambridge, 1984.

[5] Sonja Buchegger and Jean-Yves Le Boudec. Performance analysis of the CONFIDANT protocol: Cooperation of nodes – fairness in dynamic ad hoc networks. In *Proceedings of the 3rd ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pages 226 – 236, 2002.

[6] Sonja Buchegger and Jean-Yves Le Boudec. A robust reputation system for mobile ad-hoc networks. Technical Report IC/2003/50, EPFL-IC-LCA, Lausanne, Switzerland, Feb 2004.

[7] Mao Chen and Jaswinder Pal Singh. Computing and using reputations for Internet ratings. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 154–162. ACM Press, 2001.

[8] Ernesto Damiani, De Capitani di Vimercati, Stefano Paraboschi, Pierangela Samarati, and Fabio Violante. A reputation-based approach for choosing reliable resources in peer-to-peer networks. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 207–216. ACM Press, 2002.

[9] Chrysanthos Dellarocas. Analyzing the economic efficiency of eBay-like online reputation reporting mechanisms. In *Proceedings of the 3rd ACM conference on Electronic Commerce*, pages 171–179. ACM Press, 2001.

[10] eBay. http://www.ebay.com/, 2004.

[11] Jeffrey Ely, Drew Fudenberg, and David K. Levine. When is reputation bad? Discussion Paper 2035 Harvard Institute of Economic Research, Cambridge MA, May 2004.

[12] Jeffrey Ely and Juuso Välimaki. Bad reputation. *NAJ Economics*, 4(2), 2001.

[13] Epinions. http://www.epinions.com/, 2004.

[14] Michal Feldman, Kevin Lai, Ion Stoica, and John Chuang. Robust incentive techniques for peer-to-peer networks. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 102–111. ACM Press, 2004.

[15] Eric Friedman and Paul Resnick. The social cost of cheap pseudonyms. *Journal of Economics and Management Strategy*, 10(2):173–199, 2001.

[16] Drew Fudenberg and David K. Levine. Reputation and equilibrium selection in games with a patient player. *Econometrica*, 57(4):759–778, July 1989.

[17] Drew Fudenberg and Jean Tirole. *Game Theory*. MIT Press, Cambridge, Massachusetts, 1991.

[18] Andrew Gelman, John B. Carlin, Hal S. Stern, and Donald B. Rubin. *Bayesian Data Analysis*. Chapman & Hall / CRC, 2004.

[19] Gerd Gigerenzer and Reinhard Selten, editors. *Bounded rationality: The Adaptive Toolbox*. MIT Press, Cambridge, Massachusetts, 2002.

[20] Minaxi Gupta, Paul Judge, and Mostafa Ammar. A reputation system for peer-to-peer networks. In *Proceedings of the 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, pages 144–152. ACM Press, 2003.

[21] John C. Harsanyi. Games with incomplete information played by Bayesian players, part I. The basic model. *Management Science*, 14(3):159 – 182, November 1967.

[22] John C. Harsanyi. Games with incomplete information played by Bayesian players, part II. Bayesian equilibrium points. *Management Science*, 14(5):320 – 334, January 1968.

[23] John C. Harsanyi. Games with incomplete information played by Bayesian players, part III. The basic probability distribution of the game. *Management Science*, 14(7):486 – 502, March 1968.

[24] Daniel Houser and John Wooders. Reputation in Internet auctions: Theory and evidence from eBay. Working paper 00-01, University of Arizona, 2002.

[25] Bernardo A. Huberman and Fang Wu. The dynamics of reputations. *Journal of Statistical Mechanics: Theory and Experiment*, 1:P04006, April 2004.

[26] David M. Kreps and Robert Wilson. Reputation and imperfect information. *Journal of Economic Theory*, 27:253–279, 1982.

[27] David M. Kreps and Robert Wilson. Sequential equilibria. *Econometrica*, 50(4):863 – 894, July 1982.

[28] Harold W. Kuhn. Extensive games and the problem of information. In *[29]*, pages 46 – 68. 1953.

[29] Harold W. Kuhn, editor. *Classics in Game Theory*. Princeton University Press, Princeton, New Jersey, 1997.

[30] Michael W. Macy and John Skvoretz. The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63(5):638 – 660, Oct 1998.

[31] Sergio Marti and Hector Garcia-Molina. Limited reputation sharing in P2P systems. In *Proceedings of the 5th ACM Conference on Electronic Commerce*, pages 91–101. ACM Press, 2004.

[32] Geoffrey McLachlan and David Peel. *Finite Mixture Models*. Wiley-Interscience, 2000.

[33] Pietro Michiardi and Refik Molva. Core: a collaborative reputation mechanism to enforce node cooperation in mobile ad hoc networks. In *Proceedings of the IFIP TC6/TC11 Sixth Joint Working Conference on Communications and Multimedia Security: Advanced Communications and Multimedia Security*, pages 107 – 121. Kluwer, 2002.

[34] Paul Milgrom and John Roberts. Predation, reputation and entry deterrence. *Journal of Economic Theory*, 27:280–312, 1982.

[35] Nolan Miller, Paul Resnick, and Richard Zeckhauser. Eliciting informative feedback: The peer-prediction method. Working paper, Michigan School of Information and Harvard University, 2004, June 2004.

[36] Lik Mui, Mojdeh Mohtashemi, and Ari Halberstadt. Notions of reputation in multi-agents systems: a review. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 280–287. ACM Press, 2002.

[37] Richard E. Neapolitan. *Learning Bayesian Networks*. Prentice Hall, Upper Saddle River, NJ, 2004.

[38] Petteri Nurmi. Modelling routing in wireless ad hoc networks with dynamic Bayesian games. In *Proceedings of the 1st International Conference on Sensor and Communication Networks (SECON)*. IEEE, October 2004.

[39] Josep M. Pujol, Ramon Sanguesa, and Jordi Delgado. Extracting reputation in multi agent systems by means of social network topology. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 467–474. ACM Press, 2002.

[40] Paul Resnick and Richard Zeckhauser. Trust among strangers in Internet transactions: Empirical analysis of eBay's reputation system. Working paper, Michigan School of Information and Harvard University, Feb 2001.

[41] Paul Resnick, Richard Zeckhauser, John Swanson, and Kate Lockwood. The value of reputation on eBay: A controlled experiment. Working paper, Michigan School of Information, Harvard University and eBay, March 2004.

[42] Jordi Sabater and Carles Sierra. REGRET: reputation in gregarious societies. In Jörg P. Müller, Elisabeth Andre, Sandip Sen, and Claude Frasson, editors, *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 194–195, Montreal, Canada, 2001. ACM Press.

[43] Jordi Sabater and Carles Sierra. Reputation and social network analysis in multi-agent systems. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems*, pages 475–482. ACM Press, 2002.

[44] Ali Aydin Selçuk, Ersin Uzun, and Mark Reşat Pariente. A reputation-based trust management system for P2P networks. In *Proceedings of the IEEE International Symposium on Cluster Computing and the Grid (CCCGrid)*, pages 251–258. IEEE, April 2004.

[45] Reinhard Selten. The chain-store paradox. *Theory and Decision*, 9:127–159, 1978.

[46] Robert L. Trivers. The evolution of reciprocal altruism. *The Quarterly Review of Biology*, 46(1):35–57, March 1971.

[47] Yao Wang and Julita Vassileva. Bayesian network trust model in peer-to-peer networks. In *Proceedings of IEEE International Conference on Web Intelligence (WI)*, pages 372–378. IEEE, October 2003.

[48] Jorgen W. Weibull. *Evolutionary Game Theory.* MIT Press, Cambridge, Massachusetts, 1997.