Quantifying Agent Strategies Under Reputation

Sergio Marti and Hector Garcia-Molina
Stanford University
{smarti, hector}@cs.stanford.edu

Abstract

Our research proposes a simple buyer/seller game that captures the incentives dictating the interaction between peers in resource trading peer-to-peer networks. We prove that for simple reputation-based buyer strategies, a seller's decision whether to cheat or not is dependent only on the length of its transaction history, not on the particular actions committed. Given a finite number of transactions, a peer can compute a utility optimal sequence of co-operations and defections. With the limited information provided by many reputation systems, a peer has incentive to defect on a large fraction of its transactions. If temporal information is used, equilibrium is reached when peers predominantly cooperate.

I. Introduction

Peer-to-peer networks, where autonomous peers exchange resources for money or other resources, often have little or no enforcement of trading rules by a central authority. Consequently, the peers must police themselves by detecting and ostracizing agents that defect on transactions in order to mitigate such misbehavior. To detect cheaters, the P2P system employs a mechanism that tracks each peer's behavioral history so any user can determine the likelihood another party will defect on a transaction. We refer to a peer's behavioral history as its *reputation*, and the mechanism by which it is recorded, the *reputation system*. Reputation systems not only allow agents to choose more wisely who to interact with, but also discourage agents from cheating in the first place.

Reputation systems are currently used throughout the Web. Well-known examples include eBay [3], Epinions [4] and OpenRatings [15]. In addition, the burgeoning field of peer-to-peer systems research is replete with work on distributed reputation systems for networks of untrusted peers [2], [8], [9].

While most distributed reputation system work has focused on developing specific protocols and implementation designs that are tested through simulations, we believe much could be learned through high-level theoretical analysis. In particular, we take a microeconomic approach, primarily concentrating on individual transactions between a small number of buyers and sellers. In this paper, we apply game theory and economic principles to the study of reputation among autonomous peers, which we will refer to as *agents* in keeping with the source literature.

Many proposed reputation systems treat reports of success or failure equally when computing a reputation, not taking into account the time or order of successes and failures. Even eBay primarily uses the percentage of feedback that is positive as the user rating. We demonstrate how reputation systems that do not provide time information to users are vulnerable to certain adversary strategies. In fact, we prove that the optimal strategy for selfish agents gives them incentive to cheat. However, providing a more detailed transaction history to peers encourages more cooperation.

This work concentrates on selfish peers that seek to maximize their profit, regardless of the harm to other parties. We do not discuss malicious agents that gain additional utility from harming other peers, although many reputation systems are designed specifically to root out such agents [9], [12]. The analysis compares the outcome of selfish strategies to the social optimum.

For this paper, we are interested in agents that have engaged in a number of trades and therefore have accumulated a behavioral history. We ignore the issue of bootstrapping reputation for new agents while preventing whitewashing. We suggest a ”stranger adaptive” technique similar to that proposed in [5] would be effective. Also, we do not address how the behavioral history is collected. We simply assume that a perfect history is available to all agents, allowing us to focus on agent strategies rather than on specific mechanisms for gathering transaction information.

We begin in Section II by proposing a simple economic game that captures the mechanics of transactions between a buyer and a seller. In Section III, we assume a perfect reputation system and show that Nash equilibrium is reached

TABLE I. Parameter descriptions

Param.	Description
v	Valuation of $1G$ of goods to a buyer
c	Seller's production cost of $1G$ of goods
p	Price paid for a bundle
g	Fraction of bundle that is good

TABLE II. General payoff matrix

$Bundle(S)$	Buyer	Seller	Social Profit
$[1G:0B]$	$v-p$	$p-c$	$v-c$
$[0G:1B]$	$-p$	p	0
$[g:(1-g)]$	$vg-p$	$p-cg$	$(v-c)g$

when players predominantly cooperate. Then, Section IV examines equilibrium in the game given a limited number of iterations. Section V discusses related work. Finally, we conclude in Section VI.

II. Definitions and Dimensions

This section defines a game that provides a simplified model of a generic trading system. Next, we describe three dimensions which we vary to compose the specific game scenarios we are interested in analyzing.

A. Game Setup and Rules

The players in our system are buyers and sellers.

- A seller can provide 1 unit of goods each turn, which we refer to as a *bundle*. This bundle may be split by the seller between good resources, denoted by G, and bad resources, denoted by B. Let $0 \leq g \leq 1$ denote the fraction of the bundle made up of good resources. For example, bundle $[\frac{3}{4}G : \frac{1}{4}B] \Leftrightarrow g = \frac{3}{4}$.
- Each unit of good resources costs a seller c to supply and has a valuation of v to the buyer. Assume $v > c$. If not, there would be no price at which both the seller and buyer could profit from a transaction and so no transactions would occur.
- Each unit of bad resources costs a seller \$0 to supply and has a valuation of \$0 to the buyer.
- All sellers have the same production costs and all buyers have the same valuation.
- A buyer can purchase at most one bundle per turn, but may choose not to purchase any.
- The buyer always pays the seller before receiving the bundle. Consequently, the buyer can never cheat a seller, only vice versa. This assumption reduces the complexity of case analysis and mirrors most transactions, where payment is verified before goods are received and their quality evaluated.

The parameters are listed with descriptions in Table I, along with default values used in examples throughout the paper.

As with most economic games, our interest will be to study how various strategies affect the utility of each player in the game. Therefore, all values given are in units of utility. We will use \$ as the symbol for units of utility. Each player is solely motivated to increase his own utility. When a buyer purchases goods from a seller, we are interested in the change in utility for each participant of the transaction. We refer to this change in utility as the *profit* (positive or negative) of each player. We define *social profit* to be the sum of all the players' profits. We consider the optimal utilitarian strategy to be one that maximizes social profit.

B. Single Transaction Payoff

For a transaction the buyer's payoff equals the valuation of the good component of the bundle minus the price paid (p): $vg - p$. The seller receives the price minus the cost of producing the good component: $p - cg$. Adding the two gives the social profit of $(v - c)g$. These expressions are summarized in Table II for easy reference, including the two extreme bundles, $1G$ and $1B$.

These expressions hold regardless of the strategy employed by players, the number of players, or the information available to each player. Instead, these factors affect: the bundle chosen by each seller, whether a buyer agrees to buy a bundle, and the price offered by the buyers in the variable-priced scenario.

Remember, the buyer may always decline the transaction resulting in \$0 profit for both parties. If the seller is allowed to only produce $[1G : 0B]$ or $[0G : 1B]$ bundles, this game resembles the one-sided prisoner's dilemma [17], where it is one player's interest to defect when the other cooperates, while the other player wants to strictly cooperate.

C. Social Optimum

Our objective function is to maximize social profit, which we define as the sum of utility gained/lost by both the buyer and the seller. From Section II-B we have the social profit from a transaction as $(v-c)g$. Since $v-c > 0$ by definition, clearly the social optimum results when the seller maximizes g by producing $1G$.

Because social profit is independent of price or player strategy, this social optimum holds for both fixed and variable pricing and is constant across knowledge-space and player-space as well.

As we will see, the social optimum is an equilibrium for selfish agents in certain scenarios. An additional advantage of this social optimum is that it does not require the seller to know the valuation of the buyer, as long as $v > c$.

D. Dimensions

Our investigation breaks down the range of options in three dimensions: knowledge, players, and pricing. The following describes each dimension as well as the scenarios we consider relevant.

1) Knowledge-space: As we wish to look at the effects of reputation information on market behavior, we must specify what information about the seller is available to the buyer. We have looked at three approaches of increasing complexity.

Zero Knowledge (0K) Buyers have no knowledge whatsoever of the transaction history of any seller, even of sellers he himself has previously interacted with.

Perfect Knowledge (PK) Buyers know exactly what is the composition of the current bundle being offered by any seller.

Perfect History (PH) Buyers are aware of the composition of every bundle each seller has previously sold but not the current bundle.

In this paper we will focus exclusively on perfect history, which represents an ideal reputation system capable of supplying the buyer with all information about any seller's previous actions. As stated in the introduction, we use perfect history to remain implementation agnostic. However, in future work we plan to relax this assumption by incorporating probabilistic errors into seller history.

2) Player-space: The number of each type of player in a a scenario is determined as follows:

1B-1S The simplest player scenario we will look at is a game with one buyer and one seller.

1B-MS In this scenario there is one buyer but many sellers competing for the buyer's attention and money.

MB-1S Conversely, there may be many buyers competing to purchase from only one seller.

In the extended version of this paper [13], we consider more complex player scenarios, such as multiple buyers *and* sellers. In most situations each of these cases reduce to one of the three simpler scenarios, depending on relative population size. We leave out this analysis from here for brevity and clarity.

We will refer to a single buyer/seller as B/S, respectively.

3) Price-space: The two pricing options we consider are:

Fixed price (FP) The system sets a constant price for each bundle. The seller may vary the content of the bundle and the buyer may choose to buy a bundle or not, but the price does not vary.

When multiple buyers are interested in a single seller in one turn, we assume the buyers are randomly ordered. The first buyer chooses from all sellers and the rest of the buyers choose from the remaining sellers, in order. This ordering represents a real world phenomenon where an implicit ordering is obtained as buyers compete for items offered on a "first come, first served" (FCFS) basis.

Variable price (VP) Each buyer bids on a bundle offered by the seller. The seller accepts the highest bid, which determines the price the buyer must pay the seller. In the case of a tie, the seller randomly chooses.

We do not concern ourselves with the specific mechanism of the auction, but for simplicity assume an ascending auction or Vickrey auction [18].

Because auctioning bundles does not make sense when there is only one buyer we will ignore scenarios involving one buyer (i.e. 1B-1S/VP or 1B-MS/VP).

III. Basic Reputation-based Strategies

We begin by proposing very simple strategies for both buyers and sellers, then incrementally modifying them in response to the other players' current strategy until the players reach a Nash equilibrium.

As defined in Section II, perfect history (PH) entitles all buyers to know the transaction history of every seller. We will simplify our model to allow sellers to sell one of two bundles: $1G$ or $1B$. If the seller offers $1G$ we say the seller *cooperates* on the transaction. If she offers $1B$, she is *defecting* on the transaction. We argue that assuming a binary bundle does not greatly weaken our model. A buyer's decision on whether to buy, and at what price, will be based on the probability he expects the seller to cooperate or defect, which is estimated from the seller's history/reputation.

We assume that each seller has accrued a number of transactions in her history consistent with the strategy she employs. We do not focus on the reputation bootstrapping problem (when a seller has no history) which is outside the scope of this paper. When necessary we simply assume buyers expect sellers to cooperate on the first transaction.

To simplify our initial analysis of strategies for both buyers and sellers, we begin with buyers assuming a simple model for the behavior of each seller. Given this assumption, a buyer will choose a strategy. Sellers then assume buyers follow that strategy and choose their own strategy. We repeat the process, until the progression of strategies reaches Nash equilibrium where neither player has incentive to change their strategy.

This section proposes initial strategies for both buyers and sellers. The next section explores improved strategies under the auction scenario (MB-1S/VP).

Coin Model (CM): *Each round, seller S randomly chooses whether to cooperate or defect with probability ρ_S of cooperating.*

This simple model mimics each seller flipping a biased coin each turn. If there are multiple sellers in the system, each seller may have a different bias ρ_i $i \in \{S\}$ where $\{S\}$ is the set of all sellers, whether one or more.

Buyer Strategy 1 (BS1)*: Buyer B assumes seller S follows the coin model, estimates S's probability of cooperating and will pay up to $v\hat{\rho_S}$.*

Regardless of the number of buyers and sellers (*B-*S), each buyer initially considers each seller S independently. To determine the likelihood of S cooperating on the next transaction, B needs to know ρ_S. Given ρ_S the estimated valuation of S's bundle is $v\rho_S + 0(1-\rho_S) = v\rho_S$. B will be willing to pay up to $v\rho_S$ for S's bundle. Consequently, the price a seller can command is proportional to her reputation. This intuitive result is supported by empirical findings [10].

Although B may not know ρ_S, it can estimate it by using the seller's transactional history. Specifically, by counting the number of transactions it has previously cooperated on and dividing by the total number of transaction we have an unbiased estimator for ρ_S. Let T_S be the total set of transactions S participated in and C_S be those transactions in which S cooperated.

$$\hat{\rho_S} = \frac{|C_S|}{|T_S|} \tag{1}$$

To understand how a buyer uses $\hat{\rho}$ we first consider 1B-1S. Buyer B calculates $\hat{\rho_S}$ and is willing to purchase from S if the fixed price (FP) $p \leq v\hat{\rho_S}$. If S is auctioning the bundle (MB-1S/VP), B will offer at most $v\hat{\rho_S}$.

Suppose there are multiple sellers for buyer B to choose from, which is only reasonable in a fixed price scenario (1B-MS/FP). B estimates $\hat{\rho}_i$ $\forall i \in \{S\}$. A single buyer seeks to maximize expected profit $v\hat{\rho}_i - p$ for fixed price p. Thus, B will choose to purchase from S whose $\hat{\rho_S} \geq \hat{\rho}_i$ $\forall i \in S$.

While BS1 seems rather naïve, it represents the information given by many actual reputation systems, the fraction of all interactions that were positive.

In BS1 the buyer(s) assumes the seller applies the coin model. Now we will look at how the seller should respond if it assumes buyers are using BS1.

Seller Strategy 1 (SS1)*: Seller S assumes buyer uses BS1. S follows coin model, but can choose appropriate ρ_S when it enters the system. However, S cannot vary ρ_S over time.*

In other words, we allow S to freely choose ρ_S but not vary it over time (we relax this constraint in the following sections). In the 1B-1S FP scenario, ρ_S needs to be sufficiently high so that the expected valuation calculated by the buyer is greater than or equal to the price p. Therefore, $v\rho_S \geq p \Rightarrow \rho_S \geq \frac{p}{v}$. The same result holds for MB-1S FP.

However, in 1B-MS/FP, S expects the buyer B applying BS1 to choose the seller with the highest $\hat{\rho_S}$. All sellers will choose $\rho = 1$. If not, if all sellers choose a $\rho < 1$, then one seller i could unilaterally raise her ρ_i above that of the other sellers and guarantee that she is chosen by B. This move would prompt the other sellers to increase their ρ to be competitive, until all sellers are using $\rho = 1$. If one seller does not follow suit, and keeps her $\rho < 1$, then there is no chance of B choosing her.

Rational sellers will set $\rho = 1$ only if they can expect to have positive profits. Note that, if all sellers adhere to Seller Strategy 1 the expected payoff every round for a seller S is $\frac{p-c}{n}$ if $\rho_S = 1$ where n is the number of sellers with $\rho = 1$. Since only $\rho = 1$ generates positive profit for the seller, we would not expect any rational sellers to choose a $\rho < 1$.

Finally, consider the MB-1S/VP scenario. Following the same reasoning as for MB-1S VP under perfect knowledge we once again find the preferred ρ_S for seller S is 1 as long as $v > c$.

If all sellers must adhere to SS1, then the buyers have no incentive to deviate from BS1, resulting in Nash equilibrium.

IV. Independent Decisions for MB-1S/VP

In this section we consider only the one seller, multiple buyer, variable priced, perfect history scenario. The work also applies to multiple sellers, but where buyers are not restricted to purchasing at most one bundle per turn, thus allowing them to bid in each seller's auction. This scenario represents the type of markets we are most interested in, namely eBay-style auctions. Instead of insisting on a constant ρ over all time as with SS1, we allow the seller to decide whether to cooperate or defect on each transaction separately. As we will see, a crucial factor in the seller's strategy is the total number of transactions the seller plans to execute.

Suppose seller S has committed n transactions, m of which were good and $n-m$ were bad. Assuming variable priced bids and buyers applying Buyer Strategy 1, a buyer will bid up to $v\frac{m}{n}$ for the next bundle offered by seller. Should the seller cooperate or defect? If she cooperates the expected bid price of the next bundle will be $v\frac{m+1}{n+1}$. If the seller defects she gains a one-time benefit of c, but the expected price of the next bundle will be $v\frac{m}{n+1}$, slightly lower than if she had cooperated. Regardless of S's previous or successive behavior, how many additional transactions must S perform before the long-term damage done to her reputation by one defection outweighs the one time gain from that defection?

To measure the effect of a seller's decision on long-term utility we calculate utility over time for each case,

$$U_c(n+1+k) = U(n) + \Big(v\frac{m}{n} - c\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)+1}{n+i} - f(n+1+i)c\Big) \quad (2)$$

$$U_d(n+1+k) = U(n) + \Big(v\frac{m}{n}\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)}{n+i} - f(n+1+i)c\Big) \quad (3)$$

Fig. 1. Seller utility under MB-1S/VP dependent on cooperation or defection on turn $n+1$.

cooperate or defect, and see after how many rounds the values are equal.

Lemma 4.1: Assuming buyers follow BS1, a seller S that has committed n transactions will gain more utility from defecting rather than cooperating on the $n+1$ transaction if S performs less than k additional transactions and less utility if S performs more than k additional transactions, where $k \approx (n+\frac{1}{2})(e^{c/v}-1)$.

Proof: Suppose seller S has a history of n transactions, in m of which she cooperated. On the $n+1$ turn the seller chooses either to defect or cooperate. Let k be the number of turns S sells bundles after she cooperates/defects on turn $n+1$.

Let $U(n)$ be S's utility after the first n turns. Let $U_c(z)$ be the utility of S after $z > n$ turns, assuming S cooperated on the $n+1$ turn. Similarly, let $U_d(z)$ be the utility of S after $z > n$ turns, assuming S defected on the $n+1$ turn. Before we formulate $U_c(z)$ and $U_d(z)$ we must define some auxiliary functions.

Define function $f_S(t)$ to return 1 if S cooperated (C) on turn t, or 0 otherwise.

Define function $F_S(t)$ to be a nondecreasing function equal to the number of turns S has cooperated after t turns. For example, since S cooperated m times in her first n transactions, $F_S(n) = m$.

$$F_S(t) = \sum_{i=1}^{t} f_S(i) \quad (4)$$

Define $F_S^{-y}(t)$ to be a non decreasing function equal to the number of turns S has cooperated after t turns, excluding turn y. $F_S^{-y}(t) = F_S(t) - f_S(y)$. Basically, for any t the value of $F_S^{-y}(t)$ is independent of how S acted on turn y. Expressed mathematically,

$$\forall t, y\ \{F_S^{-y}(t)|f_S(y)=1\} = \{F_S^{-(y)}(t)|f_S(y)=0\} \quad (5)$$

Function $F_S^{-(n+1)}(t)$ allows us to express the fact that seller S is consistent as to whether she defects or cooperates after turn $n+1$ regardless of the decision she made that turn.

As stated above, buyers follow the consistency bid model described by BS1, therefore each turn S is paid the fraction of transactions she has cooperated in the past times the value of cooperation, v.

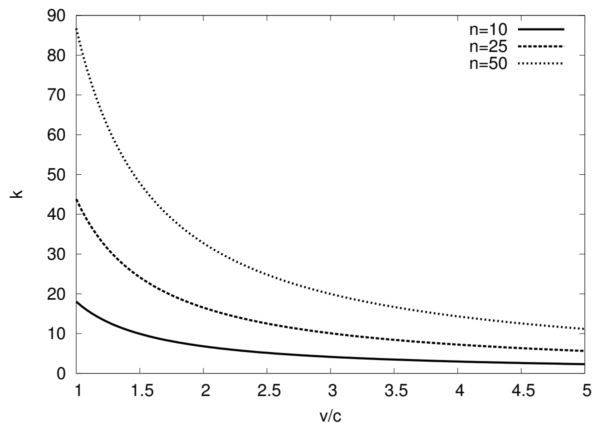

Fig. 2. Num. transactions for equalized utility after C/D decision (k) w.r.t. v/c ratio.

The seller's utility k turns after the cooperate/defect choice is expressed by the equations in Figure 1.

Setting the two utility equations equal to each other we solve for k. The full derivation is presented in the extended paper [13].

$$U_c(n+1+k) = U_d(n+1+k) \quad (6)$$

$$k \approx (n+\frac{1}{2})(e^{c/v}-1) \quad (7)$$

■

Using our default parameter values ($v=3$ and $c=1$) results in $k \approx 0.40n + 0.2$, which means a seller that has accumulated a history of 10 transactions would profit more from cooperating on the next sale than defecting if she plans to participate in 5 or more additional transactions.

Clearly there is a linear relation between n and k. For example, if $n=40$ and $v/c=2$, then $k=13.3$. Given that a buyer's valuation of a good bundle is twice the cost of producing the bundle, a seller with a history of 40 sales (good or bad) will profit less from defecting than cooperating on the next sale, if she sells 14 or more additional bundles. Interestingly, k does not depend on m or $f(x)$, only n. This means that the seller's decision to cooperate or defect on past or future transactions has no impact on whether she should cooperate or defect on the current turn; only the quantity of past and future transactions matters.

The other factor affecting k, in addition to the length of a seller's history (n), is the cost and valuation of goods. More specifically, as valuation increases with respect to cost, the optimal fraction of total transactions to defect on decreases. The ratio of cost to valuation is illustrated in Figure 2 for three values of n.

Intuitively, as the difference between cost and valuation shrinks, the potential for profit goes down. For instance, if valuation equals cost plus a small δ, then the highest price buyers will be willing to pay is the cost of the bundle plus δ. If the profit a seller can make from the sale of a good bundle is a fraction of the cost, then the utility earned by saving on the cost of one bundle outweighs the profit loss on many good bundles. This is represented by the sharp rise in k as v/c approaches 1 in the figure. As the cost of producing a good bundle becomes a smaller fraction of the valuation, and thus the bid price the seller can command for a bundle, then the decrease in bid prices due to lower reputation quickly usurps the one-time gain from defection. As v/c approaches ∞, k converges to 0.

This analysis suggests a new seller strategy for the MB-1S/VP/PH scenario:

***Seller Strategy 2**: Seller S assumes buyer uses BS1. Suppose S knows beforehand how many total bundles she wants to sell, Z, and the cost and valuation of bundles. S will maximize her utility by cooperating on the first $\lceil (Z - \frac{1}{2})e^{-c/v} - \frac{1}{2} \rceil$ transactions and then defecting on the rest.*

If S knows the total number of bundles she will auction over her lifetime in the system (call this Z), S can maximize her profit by cooperating for some of the initial transactions then, at a certain point, switching and defecting on the rest. Lemma 4.1 gives, for a certain number of completed transactions, how many more transactions must be completed for the one-time gain from defecting to equal the long-term loss due to a lower reputation. If a seller defects and performs less than k additional transactions, the defection was in her benefit. If S performs more than k, then she has less utility than had she cooperated. Therefore, ideally S's strategy is to cooperate on all sales for a number of turns, then defect on the rest of the turns. When the number of transactions in the cooperating phase is n, the number of transactions in the defecting stage is $k+1$, and the two values are related by Lemma 4.1, the utility is maximized.

Below we prove SS2 is optimal for a seller participating in a predetermined number of transactions under the scenario MB-1S/VP/PH where the buyers are using BS1.

Definition Let a *transaction schedule* of length Z be a permutation of exactly Z cooperations and defections. Let Ξ_x^Z be the set of all possible transaction schedules with x cooperations and $Z - x$ defections. For example *(C C D D C D C)*$\in \Xi_4^7$.

Definition The *utility of a transaction schedule T, $U(T)$,* is the total utility gained or lost by a seller who commits exactly Z transactions and cooperates or defects in the order specified by T, assuming MB-1S/VP with buyers using strategy BS1. $U(T)$ for any schedule T is equal to the sum of the payment received for each bundle minus the sum of the cost of producing good bundles. The total cost for a transaction $T \in \Xi_x^Z$ is xc (cooperations times cost of each). The payment received by a seller S for each bundle is equal to the buyers' valuation of a good bundle times $\hat{\rho_S}$ which, for the ith bundle, is the number of cooperations in the first $i-1$ turns divided by $i-1$. For simplicity we will assume $\hat{\rho_S} = 1$. As stated earlier, we assume buyers always trust new sellers on their first bundle. This assumption only affects the payment on the first bundle and is equal for all schedules. Mathematically,

$$\forall T \in \Xi_x^Z \quad U(T) = \underset{\substack{\text{1st} \\ \text{payment}}}{v} + \underbrace{\sum_{i=2}^{Z} v \frac{F_T(i-1)}{i-1}}_{\text{other payments}} - \underset{\text{total costs}}{x * c} \tag{8}$$

Note, the subscript in $F_T(i-1)$ refers to the transaction schedule. We define $F_T(i-1)$ as the number of cooperations in the first $i-1$ terms of transaction schedule T.

Given a seller makes Z transactions with $0 \leq x \leq Z$ cooperations and $Z - x$ defections, we will show that

(i) a utility optimal transaction schedule will consist of all x cooperations first, then all $Z-x$ defections, and
(ii) for such a transaction schedule the optimal number of cooperations is $x = \left\lceil (Z - \frac{1}{2})e^{-c/v} - \frac{1}{2} \right\rceil$.

Theorem 4.2: Assuming that buyers use strategy BS1, the utility optimal transaction schedule of length Z with x cooperations and $Z-x$ defections consists of executing all x cooperations first, followed by all $Z-x$ defections. We refer to such a schedule as a *segregated schedule*.

Proof: [Proof by contradiction] Let $\hat{T} \in \Xi_x^Z$ be an optimal transaction schedule in Ξ_x^Z such that at least one defection D appears before at least one cooperation C in the schedule. Let i be the index of the first D in $\hat{t}$ and j be the index of the last C. By definition $i < j$. Construct transaction schedule T' by swapping the D at position i with the C at position j. By definition $U(\hat{T}) \geq U(T')$. Represent each utility using Eq. 8.

$$U(\hat{T}) \geq U(T') \tag{9}$$

$$v + \sum_{k=2}^{Z} v \frac{F_{\hat{T}}(k-1)}{k-1} - x * c \geq v + \sum_{k=2}^{Z} v \frac{F_{T'}(k-1)}{k-1} - x * c \tag{10}$$

Notice both schedules have the same total cost due to having the same total number of cooperations (x). Both also have the same initial payment. Because only the i and j terms in $\hat{T}$ were swapped to form T', then $\forall k <$

$i, k \geq j \quad F_{\hat{T}}(k) = F_{T'}(k)$. Cancelling out equal terms leaves

$$\sum_{k=i+1}^{j} \frac{F_{\hat{T}}(k-1)}{k-1} \geq \sum_{k=i+1}^{j} \frac{F_{T'}(k-1)}{k-1} \qquad (11)$$

However, T' has C in the ith position where $\hat{T}$ has a D, while all other positions less than j are the same. Therefore, by definition of function $F_T(k)$, $\forall k \quad i \leq k < j \quad F_{\hat{T}}(k) = F_{T'}(k) - 1$. This fact, however, contradicts Eq. 11, which implies that $\exists k \quad i \leq k < j$ s.t. $F_{\hat{T}}(k) \geq F_{T'}(k)$. Therefore, a utility optimal transaction schedule cannot have a defection appear in the sequence before a cooperation. ■

Intuitively, because the benefit from defecting is a one-time savings on cost, while the benefit of cooperation is improved reputation which in turn increases the expected payment for each future bundle. Therefore, executing a set number of cooperations before any defections will maximize the benefit gained from those cooperations.

Theorem 4.2 implies that once a seller has decided it is in her interest to defect once, it will be in her interest to defect every time until she exits the system. Next we check to see if there is always one value of the number of cooperations that will maximize the utility of a segregated schedule of length Z.

First, we need an expression for the utility generated by a segregated schedule.

Definition Let $U_{seg}(Z, x)$ be the utility of a segregated transaction schedule of length Z with x cooperations followed by $Z - x$ defections. If we assume an MB-1S/VP scenario with buyers using BS1, $U_{seg}(Z, x)$ can be expressed as

$$U_{seg}(Z, x) = (v - c)x + \sum_{i=x}^{Z-1} v\frac{x}{i} \qquad (12)$$

where $(v - c)x$ is the utility from the x cooperations, v is the utility from the first defection, and the summation is the utility from the remaining defections. Note, as in Eq. 8, we are assuming buyers always expect the seller to cooperate on the first transaction. This assumption simplifies our derivations and analysis but does not affect our results. As we will see, for $Z \geq 2$, the seller should always cooperate on the first transaction.

Theorem 4.3: For a given value of Z the utility function for a segregated transaction schedule (given by Equation 12) has at most one unique global maximum for valid values of $0 < x \leq Z$.

Proof: The formal proof of Theorem 4.3 is given in [13]. We show that the second derivative of $U_{seg}(Z, x)$ (Eq. 12) with respect to x (the number of cooperations) is $-2v \sum_{k=0}^{\infty} \frac{k}{(x+k)^3}$, which is negative between 0 and Z. Therefore, Eq. 12 can have a most one maximum for any valid value of x. ■

Now knowing that a segregated schedule of the form *(C C . . . C D D . . . D)* with x cooperations followed by $Z - x$ defections has a unique optimal value for x that maximizes $U_{seg}(Z, x)$ for a specific Z, how can we compute it? Due to space considerations we derive an approximate answer by approximating Eq. 12 with a continuous function. We then state (Theorem 4.5) a tighter approximation based on Lemma 4.1, whose full derivation is presented in [13].

Theorem 4.4: Assuming that buyers use strategy BS1, the utility optimal transaction schedule of length Z consists of approximately $\left\lceil (Z-1)e^{-c/v} \right\rceil$ cooperations followed by $\left\lfloor (Z-1)(1 - e^{-c/v}) + 1 \right\rfloor$ defections.

Proof: Approximate $U_{seg}(Z, x)$ as the continuous function $\tilde{U}$

$$\tilde{U} = (v - c)x + \int_{x}^{Z-1} v\frac{x}{t}dt \qquad (13)$$

Simplifying and taking the derivative with respect to x yields

$$\frac{d\tilde{U}}{dx} = v \ln\Big(\frac{Z-1}{x}\Big) - c \qquad (14)$$

Set $\frac{d\tilde{U}}{dx} = 0$ and solve for x.

$$x = (Z - 1)e^{-c/v} \qquad (15)$$

Note that the second derivative of $\tilde{U}$ is

$$\frac{d^2\tilde{U}}{dx^2} = -v\frac{(Z-1)^2}{x^3} \qquad (16)$$

which is negative for all $0 < x \leq Z$. Therefore, the value of x given in Eq. 15 must give the unique maximum in $\tilde{U}$ for all valid values of x, just as for $U_{seg}(Z, x)$ (Theorem 4.3).

Because we are interested only in integer values for x and $Z - x$, the resulting equations for the optimal number of cooperations and defections in a segregated transaction schedule of length Z would be

$$n_C(Z) = \left\lceil (Z-1)e^{-c/v} \right\rceil \qquad (17)$$

$$n_D(Z) = \left\lfloor (Z-1)(1 - e^{-c/v}) + 1 \right\rfloor \qquad (18)$$

Where $n_C(Z)$ and $n_D(Z)$ are the number of cooperations and defections (respectively) in a utility optimal segregated schedule.

As stated earlier, performing a derivation based on the discrete representation of utility (presented in [13]) results in a better approximation with error bounds that approach 0 as Z approaches ∞. Restating Theorem 4.4 with a tighter approximation:

Theorem 4.5: Assuming that buyers use strategy BS1, the utility optimal transaction schedule of length Z consists of approximately $\left\lceil (Z-\frac{1}{2})e^{-c/v}-\frac{1}{2}\right\rceil$ cooperations followed by $\left\lfloor (Z-\frac{1}{2})(1-e^{-c/v})+1\right\rfloor$ defections.

We now focus solely on this improved approximation for constructing an optimal segregated schedule of length Z.

$$n_C(Z) = \left\lceil (Z-\frac{1}{2})e^{-c/v}-\frac{1}{2}\right\rceil \tag{19}$$

$$n_D(Z) = \left\lfloor (Z-\frac{1}{2})(1-e^{-c/v})+1\right\rfloor \tag{20}$$

■

Both $n_C(Z)$ and $n_D(Z)$ grow linear with respect to Z, although $n_C(Z)$ grows at a faster rate. This growth ratio is determined by the valuation/cost ratio. For our default values of c and v ($1 and $3, respectively) and a sufficiently large Z, the equations indicate a seller should cooperate on roughly the first 70% of the transactions and defect on the rest, so $n_C(Z)$ is roughly 2.5 times $n_D(Z)$.

In deriving Lemma 4.1, and consequently Eqs 19 and 20, we used a closed form approximation of a finite harmonic series. To numerically evaluate the approximation error we compute the following two error functions:

Definition From Eq. 12 we define $U_{seg}(Z, n_C(Z))$ as the utility of the transaction schedule with supposedly optimal number of cooperations and defections. Define error functions $f_e^+(Z)$ and $f_e^-(Z)$ as

$$f_e^+(Z) = \frac{U(Z,n_C(Z)) - U(Z,n_C(Z)+1)}{U(Z,n_C(Z))} \tag{21}$$

$$f_e^-(Z) = \frac{U(Z,n_C(Z)) - U(Z,n_C(Z)-1)}{U(Z,n_C(Z))} \tag{22}$$

Functions $f_e^+(Z)$ and $f_e^-(Z)$ give us the relative error between the schedule we assume to be optimal and the two closest schedules of length Z, namely with one more and one less cooperation, respectively.

The following error analysis is applied to the tighter approximation from Theorem 4.5. Due to space constraints we leave out the analogous evaluation for Theorem 4.4.

In Figure 3 we plot $f_e^+(Z)$ and $-f_e^-(Z)$. We negate the second function to better differentiate them in one graph. For large enough Z, neither curve crosses 0, indicating that indeed the schedule we believe is optimal does result in better utility than a schedule with one more or one less cooperation, and is therefore at least a local maximum. For small Z ($Z < 5$), however, $n_C(Z)$ is not necessarily optimal. In fact, though not visible in Figure 3 because it lies outside of the y-range, $f_e^+(Z)$ attains negative values for Z = 1,2,3 and 4. These results indicate that $n_C(Z) + 1$ results in better utility than $n_C(Z)$ for very small Z, which is expected because the approximation for k from Lemma 4.1 is weakest for very small Z. However, for

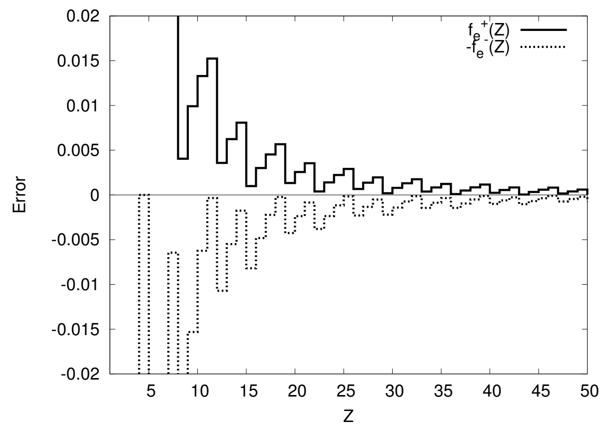

Fig. 3. Relative utility error between optimal schedule and ± 1 C/D.

very small Z the behavior of buyers towards unknown, untested sellers is an important factor. Originally, we stated we wanted to assume sufficient history in order to ignore reputation bootstrapping issues.

Calculating the same error functions using the weaker approximations from Theorem 4.4 reveals that $f_e^+(Z)$ is periodically negative, regardless of how large Z gets.

The previous numerical error analysis demonstrates that the value computed by Eq. 19 specifies the local maximum for $U_{seg}(Z, x)$. Applying Theorem 4.3, we know the value must be a global maximum because the utility function has only one unique maximum in the valid range.

Notice that if the seller plans to participate in the system selling bundles indefinitely, we may set $Z = \infty$. In this case $n_C(\infty) = \infty$. Therefore, SS2 dictates that a seller that plans to sell goods for the foreseeable future should always cooperate. As expected, this result is exactly the same as SS1, which sets $\rho = 1$.

So far we have constrained the buyers to strategy BS1. If we remove this restriction, how will buyers respond to sellers using SS2?

Buyer Strategy 2*: Buyer B assumes seller S uses SS2. Not knowing how many bundles S will sell in all (Z), B should assume S will always cooperate until S defects once. From then on assume S will always defect and never purchase from S again.*

Knowing that the optimal strategy for sellers is to cooperate for their first x transactions and then defect on the rest, a buyer will watch for a seller's first defection and then refuse to purchase any more bundles from it.

If a seller assumes all buyers are using BS2, the seller will adopt a new strategy. Knowing that no buyer will purchase a bundle from her once she has defected once, and given that the seller makes a larger profit from cooperating on a transaction than not selling anything at all, the seller will cooperate on every transaction except on the very last one.

***Seller Strategy 3**: Seller S assumes buyers use BS2. Given a total of Z bundles to sell, S will cooperate on the first $Z-1$ bundles and defect only on the last bundle.*

To this seller strategy, a buyer will respond with BS2, indicating an equilibrium. Notice SS3 is almost equivalent to SS1, where each seller cooperates on every transaction in order to maximize $\hat{\rho}$.

Using the estimator $\hat{\rho}$ alone, BS1 is unable to distinguish whether a seller, with a history of 15 Cs and 3 Ds, is applying SS1 or SS2. Obviously, a player's reputation score must rely not only on the number of cooperates and defects, but the sequence as well.

While this analysis may seem obvious, in many real world reputation systems, both proposed and deployed, only the quality ratio is often considered. For example, eBay displays the percentage of positive reviews a seller has received in her lifetime. A user can investigate a seller further and examine each and every review in chronological order. However, many users rely primarily on this single value when deciding whether to bid or not. Buyers should be better educated about evaluating the reputation information, and reputation systems that use a history weighted calculation for user quality would be more useful at discouraging defections.

V. Related Work

The work presented here was initially inspired by the work of Feldman et al. In [5], they used the Evolutionary Prisoner's Dilemma to model peer interactions in a large population. They developed a reciprocative strategy that employs subjective shared history and adaptive stranger policies to discourage selfish behavior and whitewashing. While this previous work relies primarily on simulations to evaluate the effectiveness of their design we apply mathematical analysis to derive agent strategies and overall system behavior.

Much work has applied game theory to the problem of selfish agents (e.g. [1], [6], [16]). Economists have applied game theory to market analysis and reputation for decades [11], [14], [7]. Most of this work has focused on firms competing for market share. However, the explosion in online trade among countless small transient agents demands a reevaluation of the subject. In addition, to the best of our knowledge no previous work studies optimal segregated transaction schedules.

VI. Conclusion

This paper presents a study of buyer/seller strategies, focusing primarily on how knowledge of past transaction history affects both buyer and seller strategy. We proposed a simple game model for transactions with cooperating and defecting buyers and sellers in a rich spectrum of scenarios. Beginning with basic strategies for all players, we incrementally improved them until an equilibrium was reached.

We concentrated on buyers competing in an auction (MB-1S/VP). It is interesting to note that at equilibrium players are encouraged to cooperate, realizing the social optimum. In other words, it does not pay to cheat when reputation is involved. However, buyers should be wary of the order of a seller's cooperations and defections. Maintaining only a quality ratio, as done by some reputation systems, will increase seller's incentive to cheat buyers.

References

[1] C. Buragohain, D. Agrawal, and S. Suri. A Game Theoretic Framework for Incentives in P2P Systems. In *IEEE 3rd International Conference on Peer-to-Peer Computing (P2P 2003)*.

[2] F. Cornelli, E. Damiani, and S. D. Capitani. Choosing Reputable Servents in a P2P Network. In *Proc. of the 11th International World Wide Web Conference*, 2002.

[3] eBay - The World's Online Marketplace. http://www.ebay.com/.

[4] Epinions.com. http://www.epinions.com/.

[5] M. Feldman, K. Lai, I. Stoica, and J. Chuang. Robust Incentive Techniques for Peer-to-Peer Networks. In *ACM Conference on Electronic Commerce (EC'04)*, 2004.

[6] M. Feldman, C. Padimitriou, J. Chuang, and I. Stoica. Free-Riding and Whitewashing in Peer-to-Peer Systems. In *ACM SIGCOMM 2004, Workshop of Practice and Theory of Incentives and Game Theory in Networked Systems*, 2004.

[7] D. Fudenberg and D. K. Levine. Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica*, (57), 1989.

[8] M. Gupta, P. Judge, and M. Ammar. A reputation system for peer-to-peer networks. In *ACM 13th International Workshop on Network and Operating Systems Support for Digital Audio and Video*, 2003.

[9] S. D. Kamvar, M. T. Schlosser, and H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the Twelfth International World Wide Web Conference*, 2003.

[10] C. Keser. Experimental games for the design of reputation management systems. *IBM Systems Journal*, 42(3):498–506, 2003.

[11] D. Kreps and R. Wilson. Reputation and Imperfect Information. *Journal of Economic Theory*, (50):253–79, 1982.

[12] S. Marti and H. Garcia-Molina. Limited Reputation Sharing in P2P Systems. In *ACM Conference on Electronic Commerce (EC'04)*, 2004.

[13] S. Marti and H. Garcia-Molina. A Game Theoretic Approach to Reputation (extended). Technical report, 2004. dbpubs.stanford.edu/pub/2004-49.

[14] P. Milgrom and J. Roberts. Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis. *Econometrica*, (50):443–60, 1982.

[15] OpenRatings. http://www.openratings.com/.

[16] R. Porter and Y. Shoham. Designing Efficient Online Trading Systems. In *ACM Conference on Electronic Commerce (EC'04)*, 2004.

[17] E. Rasmusen. *Games and Information*. Basil and Blackwell Ltd., 1989.

[18] W. Vickrey. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, (16):8–37, 1961.