Quantifying Agent Strategies Under Reputation

Abstract

Our research proposes a simple game model that captures the incentives dictating the interaction between buyers and sellers and reveals the strategies that evolve in different scenarios, such as eBay auctions. In particular, we find seller history has a significant effect on player strategy. 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 history, not on the particular actions committed. Given a finite number of transactions, a seller can compute a utility optimal sequence of cooperations and defections. As more advanced buyer/seller strategies evolve, equilibrium is reached when players predominantly cooperate.

1. Introduction

Online trading environments, where autonomous agents exchange resources for money or other resources, often have little or no enforcement of trading rules by a central authority. Consequently, the community of agents must police themselves by detecting and ostracizing agents that defect on transactions in order to mitigate such misbehavior. To detect cheaters, the community of agents employs a mechanism that tracks each member's behavioral history so any agent can determine the likelihood another party will defect on a transaction. We refer to an agent's behavioral history as its *reputation*, and the mechanism by which it is formed, 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 of all types are currently used throughout the Web. Well-known examples include eBay [4], Epinions [5] and OpenRatings [19]. In addition, the burgeoning field of peer-to-peer systems research is replete with work on distributed reputation systems for networks of untrusted agents [3, 9, 11].

While most 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 our previous work, we explored reputation in online trade using a macroeconomic model [17]. Now we take a microeconomic approach, primarily concentrating on individual transactions between a small number of buyers and sellers. This paper presents our initial work in applying game theory to the study of reputation.

This paper concentrates only 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 the act of harming other peers, though many reputation systems are designed specifically to root out such agents [11, 16].

For this paper, we are specifically 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 [6] would be effective. That work is further discussed in Section 7. 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 mechanisms for gathering transaction information.

We begin in Section 2 by proposing a simple economic game that captures the mechanics of transactions between a buyer and a seller. Section 3 describes the possible outcomes of each transaction and states the social optimum. In Section 4, we discuss expected player response and outcome when buyers have no knowledge about the sellers as well as when they have perfect knowledge of how a seller will respond. While simple, this exercise introduces the model and the analysis techniques, and will provide insight when we look at reputation in Section 5. Assuming a perfect reputation system, we show that Nash equilibrium is reached when players predominantly cooperate. Sections 7 and 8 discuss related and future work. Finally, we conclude in Section 10.

2. 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.

2.1. 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.

Table 1. Parameter descriptions with sample values

Param.	Description	Value
v	Valuation of $1G$ of goods to a buyer	3
c	Seller's production cost of $1G$ of goods	1
p	Price paid for a bundle	2 (for FP)
g	Fraction of bundle that is good	N/A

- 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 not want 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 1, along with default values used in concrete 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.

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.

2.2. **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 look at three approaches of increasing complexity.

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

Perfect Knowledge (PK): A buyer knows exactly what is the composition of the current bundle being offered by any seller.

Perfect History (PH): We define perfect history to mean that a buyer is aware of the composition of every bundle each seller has previously sold but not the bundle the seller is currently offering. Perfect history represents an ideal reputation system capable of supplying the buyer with all information about any seller's previous actions.

2.3. **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.

MB*MS: After studying the previous three simpler scenarios we will consider more complex player scenarios with multiple buyers and sellers, though the relative number of each will vary. The relative number of each will be indicated by the appropriate sign (i.e. =, $<$, or $>$) in place of "*". In most situations each of these cases reduce to one of the three simpler scenarios, depending on relative population size.

When the number of buyers and/or sellers does not matter we will use asterisks notation (e.g. *B-*S). We will refer to single buyers as B and single sellers as S. When there may be multiple buyers and/or sellers we will use $\{B\}$ to signify the set of all buyers and $\{S\}$ to signify the set of all sellers.

2.4. **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 [23]. Since all buyers have the same valuation for goods and the same knowledge about the seller, we expect all buyers to bid the same amount. Therefore, the second highest bid will equal the highest bid in a Vickrey auction.

Table 2. General payoff matrix

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

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).

The variable p will denote the price paid for a bundle in either price scenario. In FP, p denotes the fixed price set by the market, while in VP, p denotes the bid accepted for the bundle.

2.5. **eBay Scenario**

To help in illustrating the implications of the model, we will at times use examples within the framework of an online shopping site such as eBay [4]. While mostly known for its variable priced auctions, many items on eBay also have an associated fixed price allowing a bidder to purchase the item immediately for a specified amount. Some items are offered on a solely fixed price basis. Therefore, eBay is an excellent scenario in which to discuss the various aspects of the model across price and player space. For example, when there are more interested buyers than items offered by a particular seller at a fixed price, the order in which buyers purchase the items is determined by when each clicked the "By Now" button; first come, first served.

Though eBay covers the spectrum of player and price-space, we specially focus on the two most common scenarios: MB-1S/VP representing auctions and 1B-MS/FP representing the sale of fixed-price commodities.

3. **Strategy Independent Analysis**

This section focuses on strategy-independent properties of the model. First we discuss the payoffs each player receives from a single transaction with different bundles, as well as the social profit. Given that, we derive the socially optimal bundle.

3.1. **Single Transaction Payoff**

For a transaction the buyer's payoff equals the valuation of the good component of the bundle minus the price paid: $vg - p$. 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 2 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.

To illustrate, the following examples assume a bundle valuation of $v = \$3$, a production cost of $c = \$1$, and a fixed price of $p = \$2$ (listed in Table 1). The payoff matrix for different sample bundle distributions for these specific parameter values is given in Table 3. The last row gives the payoffs as a function of g, the fraction of the bundle that is good effort.

Table 3. Payoff matrix for fixed $\$2$ priced goods with valuation $\$3$ and cost $\$1$

$Bundle(S)$	Buyer	Seller	Social Profit
$[1G : 0B]$	1	1	2
$[\frac{1}{2}G : \frac{1}{2}B]$	-0.5	1.5	1
$[0G : 1B]$	-2	2	0
$[g : (1-g)]$	$3g - 2$	$2 - g$	$2g$

The payoff to buyers is the value of goods acquired minus the price. The payoff to the seller is the amount paid minus the cost of producing the bundle of goods. For example, consider the second row in Table 3 where a buyer purchases a bundle that is half good resources and half bad resources. The buyer gains $\frac{1}{2} \cdot \$3$ utility but pays \$2 for a total loss of \$0.5. It cost the seller \$0.5 to produce the bundle (specifically the $\frac{1}{2}G$) and it received \$2 in payment for a total gain of \$1.5. Therefore, the total increase in utility, or social profit, from the transaction was \$1.

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 [21], where it is one player's interest to defect when the other cooperates, while the other player wants to strictly cooperate.

3.2. **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 3.1 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 $v > c$.

4. Selfish Analysis

Here we compare optimal strategy previously described with player strategies due to independent selfish behavior. This section studies the 0K and PK knowledge-space while the following section focuses on the more interesting and complex Perfect History. Each part begins analyzing a one buyer-one seller scenario with fixed prices (1B-1S/FP). When applicable, variations in player-space and price-space will be discussed.

4.1. Zero Knowledge

Suppose 1B-1S/FP and consider the case of 0K, where the buyer has no knowledge of the seller's current bundle or what she has offered in the past. If every transaction is completely disconnected from all other transactions, then the seller's choice in bundle has no effect, positive or negative, on future transactions. Each round is equivalent to a one-shot Stackelberg game where the buyer always leads. Therefore, seller will offer $1B$ in order to maximize personal profit ($p - cg|g = 0 \Rightarrow p$). However, if the seller is expected to provide $1B$, purchasing from her will result in negative profit for the buyer ($vg - p|g = 0 \Rightarrow -p$). Therefore, the buyer will decline the transaction, resulting in \$0 profit for each and thus no increase in total utility.

Increasing the number of players or using variable pricing will not affect the fact that it is in each seller's interest to sell $1B$ if buyers are unable to distinguish between sellers or their bundles in any way. Therefore, it is in every buyer's interest to reject the transaction.

4.2. Perfect Knowledge

Let's begin again with 1B-1S/FP. Suppose the buyer knows exactly what bundle the seller is offering (PK). Unlike under 0K, each round is now a Stackelberg game where the seller always leads. Given that buyer B's only choices are to purchase the offered bundle or reject the transaction, seller S need only offer the minimal bundle as to give B positive profit. Solving $vg - p$ from Table 2 for $g : (1 - g)$ yields a threshold bundle of $[\frac{p}{v}G : \frac{v-p}{v}B]$. If the seller offers any bundle with more good resources, the buyer will accept. Let S offer ϵ more good resources (and thus ϵ less bad resources) where $\epsilon \to 0^+$ to ensure a very small but positive profit for B. This mixture results in a profit gain of $v\epsilon \to 0$ for the buyer and $p - c\frac{v-p}{v} - c\epsilon \to p - c\frac{v-p}{v}$ for the seller. The social profit is simply the profit of the seller, $p - c\frac{v-p}{v}$.

Using the default values for the parameters from Table 1 produces a threshold bundle of $[\frac{2}{3}G : \frac{1}{3}B]$ with a social profit of $\$\frac{4}{3}$.

Next we expand our player set to include multiple buyers, then multiple sellers.

4.2.1. *Buyer's Market.* Consider the 1B-MS/FP under PK scenario with n sellers but only one buyer. Each turn the buyer chooses the seller with the best bundle from which to purchase a bundle at the fixed price. If all sellers offer the same bundle, each seller has probability $\frac{1}{n}$ of being chosen.

Sellers can no longer offer the minimal bundle that gives a buyer a positive profit. If they do, one seller will realize that they can increase their chance of selling their bundle from $\frac{1}{n}$ to 1 by slightly improving their bundle above that of the rest. Quickly, the other sellers will follow suit and improve their bundles up to or past that of the first seller. In the end, all sellers will offer a bundle of $1G$ resulting in an average profit rate equal to the probability of being chosen times the utility gained from selling $1G$, or $\frac{1}{n} \cdot (p - c)$. No seller is motivated to change their bundle because offering anything less than the rest of the sellers guarantees they will not be chosen.

Now the Nash equilibrium equals the social optimum.

Next, consider the MB<MS scenario. Suppose there are m buyers, $m < n$. The same equilibrium will result. After $m - 1$ buyers have made a choice and chosen a seller there will remain multiple sellers for the last buyer. For these final players the problem degenerates to the 1B-MS situation and so all remaining sellers must offer $1G$. All previously chosen sellers must have also offered $1G$. If one had not, she would not have been chosen before the remaining sellers who are offering a better bundle.

4.2.2. *Seller's Market.* Next, consider the MB-1S/FP/PK scenario. Let there be m buyers and one seller. Since the seller can only sell one bundle per turn and the price of the bundle is fixed at p, she will offer the minimal bundle so as to guarantee a sale. Just as in the first case of equal number of buyers and sellers this bundle needs to be only slightly better than $[\frac{2}{3}G : \frac{1}{3}B]$ resulting in the same Nash equilibrium as $1B - 1S$.

Now we consider the variable price scenario. Instead of randomly ordering the buyers, thus guaranteeing that the last $m - 1$ will not be able to (or not want to) purchase any goods, what if we allowed the buyers to bid for bundles?

We begin with a single seller and multiple buyers (MB-1S/VP/PK). Table 4 lists the payoffs to both the buyer and the seller, as well as the total social profit for three different bundles. As expected, the social profit is the same for each bundle as in the FP scenario.

If each buyer is free to bid any price we can expect them to bid at or below the bundle valuation. Assume the valuation of the bundle is $vg > 0$ (the seller offers a bundle with at least some good content). A buyer B_1 would like to pay as little as possible, say \$0. However, a second B_2 will happily offer a bit more in order to secure winning the auction. It is in B_1's interest to raise his bid beyond that of B_2. This continues until one or both bid the actual valuation vg.

Using the MB>MS/FP scenario, we find that a similar analysis yields the same equilibrium as for the MB-1S/FP scenario. The seller will choose a bundle so as to limit the buyer's profit to 0.

Table 4. Payoff Matrix for variable priced p goods for default $v = \$3$ and $c = \$1$.

$Bundle(S)$	Buyer	Seller	Social Profit
$[1G : 0B]$	3-p	p-1	2
$[\frac{1}{2}G : \frac{1}{2}B]$	1.5-p	p-0.5	1
$[0G : 1B]$	0-p	p	0

The MB>MS/VP scenario is not as trivial. The model specifies a buyer can acquire one bundle per round. With multiple sellers auctioning their bundles should buyers be allowed to bid on multiple concurrent bundles? One solution is to order the sellers and conduct the auctions sequentially. The buyer who wins the bundle cannot participate in progressive auctions. With more buyers than sellers, we are guaranteed to have multiple bidders for each bundle and so the same equilibrium price as the MB-1S is expected.

Similarly, if we allow buyers to purchase multiple bundles, and the valuation of each bundle is not affected by the number acquired, then we would expect every buyer to participate in each seller's auction. Once again, this situation degenerates to the MB-1S case.

The last scenario is auctions held in parallel and buyers can only purchase one bundle and therefore only bid on one bundle. Here we break it down into two cases: one with more than twice as many buyers as sellers, and one with less.

The first case is the simplest. Each seller's auction will be bid on by two or more buyers, mirroring the MB-1S situation. If not, if there were a seller with only one buyer bidding on its bundle then that buyer would have an advantage and would bid low (less than v). However, there must be a seller with three or more buyers bidding for her bundle. One of those buyers would see that the single buyer was bidding less than v and move its bid over the the single-buyer seller and escalate the bid. Now every seller has multiple buyers bidding.

In the second case, there are less than two buyers for each seller. Some sellers will have only one buyer bidding for their bundles.

To summarize, the 0K results indicate the need for some information about a seller's behavior if any trades are to happen. Even with perfect knowledge, the seller will not necessarily act in the best interest of the buyer. However, in many scenarios the seller has incentive to offer the best possible bundle. While obvious in situations where multiple buyers are competing for one buyer's attention (and money), it also holds when multiple buyers are competing for one seller's item in an auction scenario.

5. **Perfect History**

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 2, 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. This probability will be estimated based on all sellers' 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. If sellers then assume each buyer follows that strategy, they will choose their own strategy. We then repeat the process, until the progression of strategies reaches Nash equilibrium where neither player has incentive to change their strategy.

The first section proposes initial strategies for both buyers and sellers. The following section explores improved strategies under the auction scenario (MB-1S/VP), while the final section concentrates on strategies in the fixed-price market (1B-MS/FP) scenario.

5.1. **Basic Reputation-based Strategies**

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 $\boldsymbol{\beta_1}$ (BS-$\boldsymbol{\beta_1}$)*: 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 $v\rho_S + 0(1 - \rho_S) = v\rho_S$. Therefore, 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 [12].

Though 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 $\hat{\rho}$ affects the buyer's decision, 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 $v\hat{\rho_S}$.

Now suppose there are multiple sellers to choose from (1B-MS). B estimates $\hat{\rho}_i \; \forall i \in \{S\}$. Now consider the following cases.

- **FP**: B seeks to maximize expected profit $v\hat{\rho}_i - p$ for fixed price p. Therefore a single buyer (1B-MS) will choose to purchase from S whose $\hat{\rho_S} \geq \hat{\rho}_i \; \forall i \in S$. If there are multiple buyers (MB*MS) competing for the bundles on FCFS, from among the remaining sellers with available bundles, B will purchase from the seller with the highest $\hat{\rho}_i$ such that $v\hat{\rho}_i - p \geq 0$.
- **VP**: Variable pricing only applies to $MB > MS$, or $MB < MS$ if buyers can purchase multiple bundles per turn. In either case, B will bid up to $v\hat{\rho_S}$, just as in the single seller scenario.

In BS-β_1 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 BS-β_1.

***Seller Strategy $\boldsymbol{\sigma_1}$ (SS-$\boldsymbol{\sigma_1}$)**: Seller S assumes buyer uses BS-β_1. 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 BS-β_1 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 σ_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 SS-σ_1, then the buyers have no incentive to deviate from BS-β_1, resulting in Nash equilibrium.

5.2. 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 SS-σ_1, 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 β_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 (compare $1G$ with $1B$ in Table 4), 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, cooperate or defect, and see after how many rounds the values are equal.

Lemma 5.1. *Assuming buyers follow BS-β_1, 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, $F_S(n) = m$.

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

$$U_c(n+1+k) = U(n) + \Big(v\frac{m}{n} - c\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)+1}{n+i} - f(n+1+i)c\Big) \tag{6}$$

$$U_d(n+1+k) = U(n) + \Big(v\frac{m}{n}\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)}{n+i} - f(n+1+i)c\Big) \tag{7}$$

Figure 1. Equations for utility of a seller under MB-1S/VP depending on whether she cooperates or defects on turn $n+1$.

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:

$$F_S^{-y}(t) = \sum_{i=1, i\neq y}^{t} f_S(i) \tag{3}$$

Basically, for any t the value of $F_S^{-y}(t)$ is independent of how S acted on turn y. Expressed mathematically,

$$\forall t, y\ \{F_S^{-y}(t) | f_S(y) = 1\} = \{F_S^{-(y)}(t) | f_S(y) = 0\} \tag{4}$$

Specifically of interest to our problem substitute y with $n+1$.

$$\begin{aligned} \forall t \quad & \{F_S^{-(n+1)}(t) | f_S(n+1) = 1\} \\ =& \{F_S^{-(n+1)}(x) | f_S(n+1) = 0\} \end{aligned} \tag{5}$$

Function $F_S^{-(n+1)}(t)$ allows us to express the fact that the seller is consistent as to whether she defects or cooperates after turn $n+1$ regardless of the decision she made that turn. As we are dealing with only one seller we will ignore the subscript.

As stated above, buyers follow the consistency bid model described by BS-β_1, therefore each turn S is paid the fraction of transactions she has cooperated in the past times the value of cooperation, v.

The following equations express the seller's utility k turns after the cooperate/defect choice is expressed by the equations in Figure 5.2.

Notice in Equation 6 the additional 1 in the summation fraction indicating S cooperated on the $n+1$ turn. Set the two utility equations equal to each other and solve for k. The full derivation is presented in Appendix A.

$$U_c(n+1+k) = U_d(n+1+k) \tag{8}$$

$$k \approx (n + \frac{1}{2})(e^{c/v} - 1) \tag{9}$$

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.

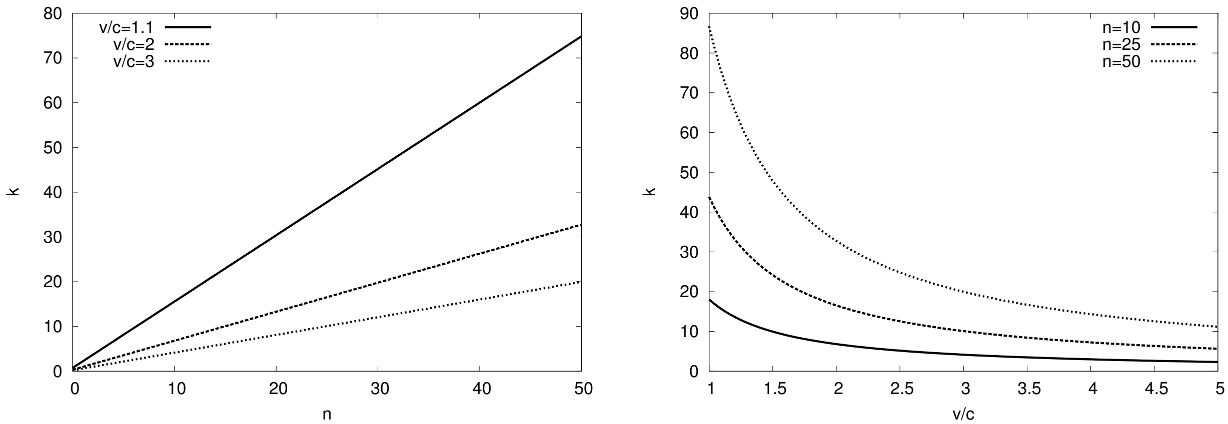

(a) As a function of seller's history n, for three values of v/c.

(b) As a function of the ratio of valuation to cost v/c, for three values of n.

Figure 2. Number of transactions until gain from single defection equals loss from lowered reputation k.

Figure 2(a) shows the linear relation between n and k for three different values of v/c. For example, for $n = 40$ and $v/c = 2$, $k = 13.3$, therefore, 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 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(b) 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 $\boldsymbol{\sigma_2}$**: Seller S assumes buyer uses BS-$\boldsymbol{\beta_1}$. 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 5.1 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 5.1, the utility is maximized.

Below we prove SS-σ_2 is optimal for a seller participating in a predetermined number of transactions under the scenario MB-1S/VP/PH where the buyers are using BS-σ_1.

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 BS-β_1. $U(T)$ 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,

$$\forall T \in \Xi_x^Z \quad U(T) = \underset{\substack{\text{1st}\\\text{payment}}}{v} + \underbrace{\sum_{i=2}^{Z} v \frac{F_T(i-1)}{i-1}}_{\text{other payments}} - \underset{\text{total costs}}{x * c} \tag{10}$$

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 5.2. *Assuming that buyers use strategy BS-β_1, 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 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. 10.

$$U(\hat{T}) \geq U(T') \tag{11}$$

$$v + \sum_{k=2}^{Z} v \frac{F_{\hat{T}}(k-1)}{k-1} - x * c \geq v + \sum_{k=2}^{Z} v \frac{F_{T'}(k-1)}{k-1} - x * c \tag{12}$$

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

$$\cancel{v} + \cancel{v} \sum_{k=2}^{Z} \frac{F_{\hat{T}}(k-1)}{k-1} - \cancel{x * c} \geq \cancel{v} + \cancel{v} \sum_{k=2}^{Z} \frac{F_{T'}(k-1)}{k-1} - \cancel{x * c} \tag{13}$$

$$\sum_{k=i+1}^{j} \frac{F_{\hat{T}}(k-1)}{k-1} \geq \sum_{k=i+1}^{j} \frac{F_{T'}(k-1)}{k-1} \tag{14}$$

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. 14, 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 5.2 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 BS-β_1, $U_{seg}(Z, x)$ can be expressed as

$$U_{seg}(Z, x) = (v - c)x + \sum_{i=x}^{Z-1} v \frac{x}{i} \tag{15}$$

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. 10, 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 5.3. *For a given value of Z the utility function for a segregated transaction schedule (given by Equation 15) has at most one unique global maximum for valid values of $0 < x \leq Z$.*

Proof The formal proof of Theorem 5.3 is given in Appendix B. Basically, the second derivative of $U_{seg}(Z, x)$ (Eq. 15) with respect to x (the number of cooperations) is $-2v \sum_{k=0}^{\infty} \frac{k}{(x+k)^3}$, which is always negative between 0 and Z. Therefore, Eq. 15 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. 15 with a continuous function. We then state (Theorem 5.5) a tighter approximation based on Lemma 5.1 and its proof in Appendix A, whose full derivation is presented in Appendix C.

Theorem 5.4. *Assuming that buyers use strategy BS-β_1, 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 $\tilde{U}$

$$\tilde{U} = (v-c)x + \int_{x}^{Z-1} v\frac{x}{t}dt \tag{16}$$

Simplifying and taking the derivative with respect to x yields

$$\tilde{U} = (v-c)x + vx\ln(Z-1) - vx\ln(x) \tag{17}$$

$$\frac{d\tilde{U}}{dx} = (v-c) + v\ln(Z-1) - v\ln(x) - v \tag{18}$$

$$= v\ln\left(\frac{Z-1}{x}\right) - c \tag{19}$$

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

$$\frac{d\tilde{U}}{dx} = v\ln\left(\frac{Z-1}{x}\right) - c = 0 \tag{20}$$

$$\ln\left(\frac{Z-1}{x}\right) = \frac{c}{v} \tag{21}$$

$$\frac{Z-1}{x} = e^{c/v} \tag{22}$$

$$x = (Z-1)e^{-c/v} \tag{23}$$

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

$$\frac{d^2\tilde{U}}{dx^2} = -v\frac{(Z-1)^2}{x^3} \tag{24}$$

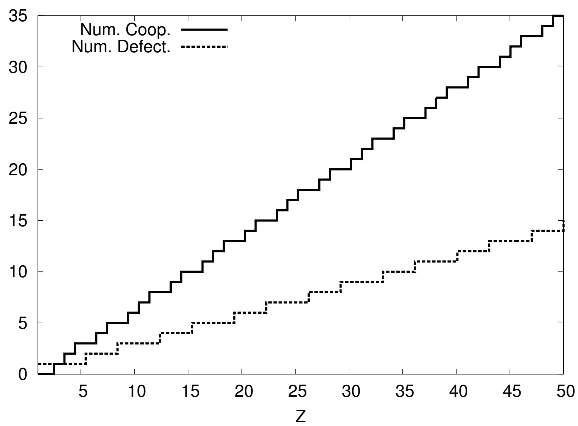

Figure 3. Optimal number of cooperation/defections as a function of total sales.

which is negative for all $0 < x \leq Z$. Therefore, the value of x given in Eq. 23 must give the unique maximum in $\tilde{U}$ for all valid values of x, just as for $U_{seg}(Z, x)$ (Theorem 5.3).

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

$$n_C(Z) = \left\lceil (Z-1)e^{-c/v} \right\rceil \tag{25}$$

$$n_D(Z) = \left\lfloor (Z-1)(1 - e^{-c/v}) + 1 \right\rfloor \tag{26}$$

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 Appendix C) results in a better approximation with error bounds that approach 0 as Z approaches ∞. Restating Theorem 5.4 with a tighter approximation:

Theorem 5.5. *Assuming that buyers use strategy BS-β_1, 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.

$$n_C(Z) = \left\lceil (Z - \frac{1}{2})e^{-c/v} - \frac{1}{2} \right\rceil \tag{27}$$

$$n_D(Z) = \left\lfloor (Z - \frac{1}{2})(1 - e^{-c/v}) + 1 \right\rfloor \tag{28}$$

Figure 3 shows both $n_C(Z)$ and $n_D(Z)$ as functions of Z for $v/c = 3$. Notice that both are linear in Z, though $n_C(Z)$ grows at a faster rate, so that it is always roughly 2.5 times $n_D(Z)$. This 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

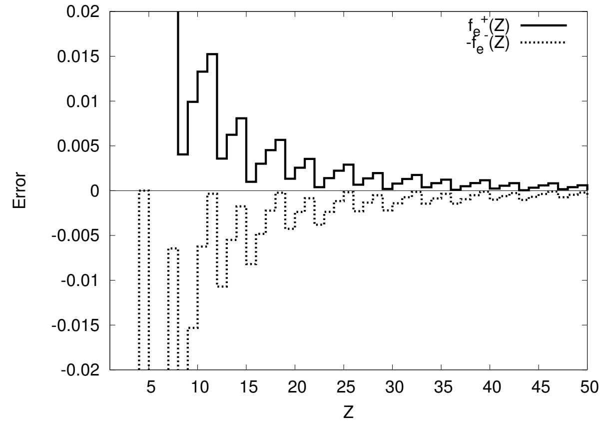

Figure 4. Relative utility error between optimal schedule and $\pm 1\ C/D$**.**

equations indicate a seller should cooperate on roughly the first 70% of the transactions and defect on the rest. For example, we see that at $Z = 40$, $n_C(Z) = 28$ and $n_D(Z) = 12$; $\frac{28}{40} = 70\%$.

In deriving Lemma 5.1, and consequently Eqs 27 and 28, we used a closed form approximation of a finite harmonic series (see Appendix A). To numerically evaluate the approximation error we compute the following two error functions:

Definition From Eq. 15 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

$$f_e^+(Z) = \frac{U(Z, n_C(Z)) - U(Z, n_C(Z)+1)}{U(Z, n_C(Z))} \tag{29}$$

$$f_e^-(Z) = \frac{U(Z, n_C(Z)) - U(Z, n_C(Z)-1)}{U(Z, n_C(Z))} \tag{30}$$

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 5.5. Due to space constraints we leave out the analogous evaluation for Theorem 5.4.

In Figure 4 we plot $f_e^+(Z)$ and $-f_e^-(Z)$.[1] 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 4 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 5.1 is weakest for very small Z. However, for 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.

[1] We negate the second function to better differentiate both functions in one graph.

Calculating the same error functions using the weaker approximations from Theorem 5.4 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. 27 specifies the local maximum for $U_{seg}(Z, x)$. Applying Theorem 5.3, 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, SS-σ_2 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 SS-σ_1, which sets $\rho = 1$.

So far we have constrained the buyers to strategy BS-β_1. If we remove this restriction, how will buyers respond to sellers using SS-σ_2?

Buyer Strategy $\boldsymbol{\beta_2}$*: Buyer B assumes seller S uses SS-σ_2. 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 BS-β_2, 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 $\boldsymbol{\sigma_3}$*: Seller S assumes buyers use BS-β_2. 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 BS-β_2, indicating an equilibrium. Notice SS-σ_3 is almost equivalent to SS-σ_1, where each seller cooperates on every transaction in order to maximize $\hat{\rho}$.

Using the estimator $\hat{\rho}$ alone, BS-β_1 is unable to distinguish whether a seller, with a history of 15 Cs and 3 Ds, is applying SS-σ_1 or SS-σ_2. Obviously, a player's reputation score must rely not only on the number of cooperates and defects, but the sequence as well.

5.3. **Independent Decisions for 1B-MS/FP**

We continue studying the expected player behavior when sellers decide on a per turn basis whether to cooperate or defect. While the previous section dealt specifically with the MB-1S/VP scenario, in this section we concentrate on the 1B-MS/FP scenario.

A rational seller's decision whether to cooperate or defect is not fixed over time (as in SS-σ_1); it may vary as both her and other sellers' reputations vary. For example, suppose there are 10 sellers, $S_1 ... S_{10}$ and one buyer B

that is willing to pay a fixed price p for one bundle and follows Buyer Strategy β_1. Each seller has previously sold 10 bundles, of which 5 were good and 5 bad. All else being equal, the buyer will prefer to purchase from the seller with the best transactional record, expressed as the fraction of transactions in which they cooperated. To the buyer who must choose one, all ten are identical and thus all have an equal chance of being chosen, 0.1. Suppose B chooses S_1. If S_1 defects, it's transaction record will drop to $\frac{5}{11}$ while the rest remain at $\frac{5}{10}$. In the following round, having a clearly worse record will disqualify S_1 from selection, lowering S_1's probability of being chosen to 0 and raising the other seller's chance to $\frac{1}{9}$. However, if S_1 had instead cooperated with B then her record would be $\frac{6}{11}$, higher than the other sellers. In the following round we would expect B to choose S_1 with probability 1 as she clearly has a better record than the rest. These expected outcomes provide incentive for S_1 to cooperate.

In the following round, B chooses S_1 again. If S_1 defects, her record falls to $\frac{6}{12}$, equal to that of the other sellers. S_1 is no longer guaranteed to be chosen and once again has a probability of 0.1. Therefore, once again, S_1 is incentivized to cooperate.

The third round, however, the situation is more interesting. If S_1 is chosen and defects her record drops to $\frac{7}{13}$, which is still better than the rest of the sellers with $\frac{5}{10}$. There is no disincentive for S_1 to defect. In fact, comparing $1G$ to $1B$ in Table 3, S_1 clearly has incentive to defect and earn more utility than cooperating.

This analysis suggests a new strategy for the seller.

Suppose there are n sellers, each with c_i cooperations and d_i defections (c_i/d_i not necessarily equal to c_j/d_j, if $i \neq j$). Assume a buyer ranks sellers according to their reputation score, calculated as $\frac{c_i}{c_i+d_i}$ for seller i. Then a seller i would choose to defect on a transaction if $\frac{c_i}{c_i+d_i+1} > \frac{c_j}{c_j+d_j}\ \forall j \neq i$. Otherwise, the seller cooperates. More generally stated,

***Seller Strategy $\boldsymbol{\sigma_4}$**: Assuming buyers use BS-β_1, seller S always cooperate unless her reputation is sufficiently higher than the other sellers that a defection still gives S a higher reputation than the rest.*

Now we relax the assumption that buyers strictly use BS-β_1.

In the example above, we noted that S_1 had incentive to defect in round 3 while keeping her standing as the highest reputable seller. Consequently, B may be better off ignoring S_1 and choosing one of the other sellers, contrary to BS-β_1. Now, the situation is reduced to the problem of 9 equal sellers, all with incentive to cooperate. Suppose B chooses S_2 now. As before S_2 is expected to cooperate, raising her reputation to $\frac{6}{11}$. On the fourth round, B will choose S_1 again, since if she defects, S_1's new reputation of $\frac{7}{13}$ will be lower than S_2's.

If sellers are expected to use SS-σ_4, a buyer should then choose the seller S, such that $\frac{c_S}{c_S+d_S} \geq \frac{c_j}{c_j+d_j}\ \forall j$ unless $\frac{c_S}{c_S+d_S+1} > \frac{c_j}{c_j+d_j}\ \forall j \neq S$. In such a case, the buyer should choose the seller T such that $\frac{c_T}{c_T+d_T} \geq \frac{c_j}{c_j+d_j}\ \forall j \neq S$. In other words,

***Buyer Strategy $\boldsymbol{\beta_3}$**: Assuming sellers use SS-σ_4, buyer B always chooses to buy from the seller with the highest reputation whose rank (from most reputable to least reputable) would fall if they defected.*

We believe that knowing buyers are using BS-β_3 will not cause sellers to deviate from SS-σ_4, and hence equilibrium is reached. We do not present a formal proof here.

6. Variably-valuated goods

So far we have discussed situations where every good sold by a seller had an equal cost. In real markets a seller sells goods of varying value. How does this affect reputation? How important is the value of the transactions in a seller's history? If no import is placed on the transaction value a seller could accumulate a high reputation selling inexpensive goods and then defect on one large transaction.

One improvement may be to use the price of each transaction when computing a seller's reputation. For example, let us assume a buyer is using BS-β_1 and wants to calculate $\hat{\rho_S}$ for seller S. Instead of using the formula in Equation 1 where each previous transaction is reduced to 0 or 1, we can use the following equation

$$\hat{\rho_S} = \frac{\sum_{i \in C_S} p(i)}{\sum_{i \in T_S} p(i)} \tag{31}$$

This estimator allows a buyer to better detect a seller who purposefully cooperates on small transactions, but defects on very large ones, and distinguish that seller from one who makes accidental errors that tarnish its reputation.

7. Related Work

The work presented here was initially inspired by the work of Feldman et al. In [6], 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, 7, 20]). [1] predicts a socially beneficial Nash equilibrium given some incentive scheme, while [7] concentrates on minimizing whitewashing. However, most of this research uses one-shot games to model behavior and do not address peer history or reputation.

In both [2] and [12], user groups participated in economic games in order to experimentally compare the market efficiency from varying amounts of transaction history. Their results are similar to our analytic results, indicating that the more information available about an agent, the more likely they will cooperate.

Economists have applied game theory to market analysis and reputation for decades [14, 18, 8]. 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.

8. Future Directions

The work presented here assumed all buyers had an equal constant valuation that was public. One extension will be to allow buyers to have different, private bundle valuations. This would only affect seller strategies that rely on knowing v in order to choose the proper course of action.

Much of the analysis of this simplified model indicated equilibrium in some scenarios is reached when sellers only cooperate. Introducing an unavoidable error rate that results in occasional defections regardless of the seller's intention, may require buyers and sellers to devise more interesting strategies. What is sellers could gain a cost reduction on all bundles by accepting a higher error rate? This would mimic retailers choosing to stock cheaper items from lower quality manufacturers.

A necessary step will be to forego our assumption of perfect history and explore the use of uncertain history provided by imperfect reputation systems. One solution would be to assign a probability that any given transaction is incorrectly reported or simply omitted from a seller's recorded history.

9. Acknowledgements

The authors would like to thank Kevin Lai and Tim Roughgarden for their comments and insights.

10. Conclusion

This paper presents our initial 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 the two scenarios we believe to be the most natural, buyers competing in an auction (MB-1S/VP) and many sellers competing for buyers in a fixed-price commodities market (1B-MS/FP). 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.

References

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

[2] Chen, K.-Y., Hogg, T., and Wozny, N. Experimental Study of Market Reputation Mechanisms. In *ACM Conference on Electronic Commerce (EC'04)* (2004).

[3] CORNELLI, F., DAMIANI, E., AND CAPITANI, S. D. Choosing Reputable Servents in a P2P Network. In *Proc. of the 11th International World Wide Web Conference* (2002).
[4] eBay - The World's Online Marketplace. http://www.ebay.com/.
[5] Epinions.com. http://www.epinions.com/.
[6] FELDMAN, M., LAI, K., STOICA, I., AND CHUANG, J. Robust Incentive Techniques for Peer-to-Peer Networks. In *ACM Conference on Electronic Commerce (EC'04)* (2004).
[7] FELDMAN, M., PADIMITRIOU, C., CHUANG, J., AND STOICA, I. 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).
[8] FUDENBERG, D., AND LEVINE, D. K. Reputation and Equilibrium Selection in Games with a Patient Player. *Econometrica*, 57 (1989).
[9] GUPTA, M., JUDGE, P., AND AMMAR, M. 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).
[10] HAVIL, J. *Gamma: Exploring Euler's Constant*. Princeton University Press, 2003.
[11] KAMVAR, S. D., SCHLOSSER, M. T., AND GARCIA-MOLINA, H. The EigenTrust Algorithm for Reputation Management in P2P Networks. In *Proceedings of the Twelfth International World Wide Web Conference* (2003).
[12] KESER, C. Experimental games for the design of reputation management systems. *IBM Systems Journal 42*, 3 (2003), 498–506.
[13] KNUTH, D. E. *Fundamental Algorithms*, 2nd ed., vol. 1 of *The Art of Computer Programming*. Addison-Wesley Publishing Co., 1973.
[14] KREPS, D., AND WILSON, R. Reputation and Imperfect Information. *Journal of Economic Theory*, 50 (1982), 253–79.
[15] LAI, K. Personal communication, 2004.
[16] MARTI, S., AND GARCIA-MOLINA, H. Limited Reputation Sharing in P2P Systems. In *ACM Conference on Electronic Commerce (EC'04)* (2004).
[17] MARTI, S., AND GARCIA-MOLINA, H. Modeling Reputation and Incentives in Online Trade (extended). Tech. rep., 2004. dbpubs.stanford.edu/pub/2004-45.
[18] MILGROM, P., AND ROBERTS, J. Limit Pricing and Entry Under Incomplete Information: An Equilibrium Analysis. *Econometrica*, 50 (1982), 443–60.
[19] OpenRatings. http://www.openratings.com/.
[20] PORTER, R., AND SHOHAM, Y. Designing Efficient Online Trading Systems. In *ACM Conference on Electronic Commerce (EC'04)* (2004).
[21] RASMUSEN, E. *Games and Information*. Basil and Blackwell Ltd., 1989.
[22] ROUGHGARDEN, T. Personal communication, 2004.
[23] VICKREY, W. Counter speculation, auctions, and competitive sealed tenders. *Journal of Finance*, 16 (1961), 8–37.
[24] WEISSTEIN, E. W. Harmonic number. From MathWorld–A Wolfram Web Resource, 2004. http://mathworld.wolfram.com/HarmonicNumber.html.
[25] WEISSTEIN, E. W. Polygamma function. From MathWorld–A Wolfram Web Resource, 2004. http://mathworld.wolfram.com/PolygammaFunction.html.

Appendix A. Proof Of Long-Term Reputation Damage

Set the two utility equations (Eq. 6 and Eq. 7) equal to each other and solve for k.

$$U_c(n+1+k) = U_d(n+1+k) \tag{32}$$

$$\begin{gathered} U(n) + \Big(v\frac{m}{n} - c\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)+1}{n+i} - f(n+i+1)c\Big) \\ = U(n) + \Big(v\frac{m}{n}\Big) + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)}{n+i} - f(n+i+1)c\Big) \end{gathered} \tag{33}$$

$$\begin{gathered} \cancel{U(n)} + \cancel{v\frac{m}{n}} - c + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)+1}{n+i}\Big) - \cancel{\sum_{i=1}^{k} f(n+i+1)c} \\ = \cancel{U(n)} + \cancel{v\frac{m}{n}} + \sum_{i=1}^{k}\Big(v\frac{F^{-(n+1)}(n+i)}{n+i}\Big) - \cancel{\sum_{i=1}^{k} f(n+i+1)c} \end{gathered} \tag{34}$$

$$-c + \cancel{\sum_{i=1}^{k} v\frac{F^{-(n+1)}(n+i)}{n+i}} + \sum_{i=1}^{k} v\frac{1}{n+i} = \cancel{\sum_{i=1}^{k} v\frac{F^{-(n+1)}(n+i)}{n+i}} \tag{35}$$

$$v\sum_{i=1}^{k}\frac{1}{n+i} = c \tag{36}$$

$$\sum_{i=1}^{n+k}\frac{1}{i} - \sum_{i=1}^{n}\frac{1}{i} = \frac{c}{v} \tag{37}$$

$$\tag{38}$$

Now we have two finite harmonic sums. To simplify the summations, we apply the formula for finite harmonic sum [13].

$$\begin{gathered} H_n = \sum_{i=1}^{n}\frac{1}{i} = \ln(n) + \gamma + \frac{1}{2n} - \frac{1}{12n^2} + \frac{1}{120n^4} - \epsilon \\ \text{where } 0 < \epsilon < \frac{1}{252n^6} \end{gathered} \tag{39}$$

Where γ is the Euler-Mascheroni constant.

$$\text{Let} \quad \epsilon'(n) = \frac{1}{2n} \tag{40}$$

$$\text{Clearly,} \quad \ln(n) + \gamma < H_n < \ln(n) + \gamma + \epsilon'(n) \tag{41}$$

Next substitute the appropriate upper or lower bound for H_n for each summation in Eq. 38 so as to get an upper and lower bound k.

$$(\ln(n+k)+\gamma) - (\ln(n)+\gamma+\epsilon'(n)) < \frac{c}{v} < (\ln(n+k)+\gamma+\epsilon'(n+k)) - (\ln(n)+\gamma) \tag{42}$$

$$\ln(n+k)+\cancel{\gamma}-\ln(n)-\cancel{\gamma}-\epsilon'(n) < \frac{c}{v} < \ln(n+k)+\epsilon'(n+k)+\cancel{\gamma}-\ln(n)-\cancel{\gamma} \tag{43}$$

$$\ln\Big(\frac{n+k}{n}\Big) - \epsilon'(n) < \frac{c}{v} < \ln\Big(\frac{n+k}{n}\Big) + \epsilon'(n+k) \tag{44}$$

(45)

Notice that $\epsilon'(n+k) \leq \epsilon'(n)\ \forall k \geq 0$. We can replace $\epsilon'(n+k)$ with $\epsilon'(n)$ without invalidating the inequality.

$$\ln\Big(\frac{n+k}{n}\Big) - \epsilon'(n) < \frac{c}{v} < \ln\Big(\frac{n+k}{n}\Big) + \epsilon'(n) \tag{46}$$

$$\frac{n+k}{n}e^{-\epsilon'(n)} < e^{c/v} < \frac{n+k}{n}e^{\epsilon'(n)} \tag{47}$$

$$(n+k)e^{-\epsilon'(n)} < ne^{c/v} < (n+k)e^{\epsilon'(n)} \tag{48}$$

(49)

Solving each inequality separately, we have

$$(n+k)e^{-\epsilon'(n)} < ne^{c/v} \qquad ne^{c/v} < (n+k)e^{\epsilon'(n)} \tag{50}$$

$$n+k < ne^{c/v}e^{\epsilon'(n)} \qquad n+k > ne^{c/v}e^{-\epsilon'(n)} \tag{51}$$

$$k < n(e^{c/v}e^{\epsilon'(n)}-1) \qquad k > n(e^{c/v}e^{-\epsilon'(n)}-1) \tag{52}$$

Notice that $e^{-\epsilon'(n)} < 1$ and $e^{\epsilon'(n)} > 1$. Therefore, we will approximate k to be

$$k \approx n(e^{c/v}-1) \tag{53}$$

A.1. **Error Bounds**

What is the error range for k? Subtracting the lower bound from the upper bound we have

$$n(e^{c/v}e^{\epsilon'(n)}-1) - n(e^{c/v}e^{-\epsilon'(n)}-1) = e^{c/v}n(e^{\epsilon'(n)} - e^{-\epsilon'(n)}) \tag{54}$$

$$\text{Consider, } \lim_{n\to\infty} n(e^{\epsilon'(n)} - e^{-\epsilon'(n)}) = \lim_{n\to\infty} n(e^{\frac{1}{2n}} - e^{-\frac{1}{2n}}) \tag{55}$$

$$\text{Substituting } x = \frac{1}{2n} \text{ gives } \lim_{x\to 0} \frac{e^x - e^{-x}}{2x} \tag{56}$$

$$\lim_{x\to 0} \frac{e^x - e^{-x}}{2x} = 1 \tag{57}$$

As $n \to \infty$ the error range decreases and converges to $e^{c/v}$, a constant with respect to n. Therefore, the largest error is when n is as small as possible. Originally, we stated that we are not concerned with the situation that the seller is new to the system, but has instead generated a history of several transactions. Therefore, we assume n is not small and definitely not 0. For example, using $n = 5$ in Eq. 54, the error range will be $1.002e^{c/v}$.

By definition $c < v$, therefore $e^{c/v} < e$. In our running example of $c = \$1$ and $v = \$3$, then $e^{1/3} = 1.4$. Here, the error range is less than 1.5. Because we are interested in k as an integer, the approximate value from Eq. 53 cannot be off by more than 1. Even in the worst case, where $n = 1$ and $e^{c/v} = e$, the error range is less than 3, therefore the approximate value of k cannot be off by more than 2. For the range of n we are interested in, the approximation of k may be acceptable.

A.2. **Improved Approximation**

In the previous section we bounded the error to a range of size $e^{c/v}$, constant with respect to n. However, it is possible to do better by using a tighter bound on the harmonic number [10, 24].

$$\frac{1}{24(n+1)^2} < H_n - \ln(n + \frac{1}{2}) - \gamma < \frac{1}{24n^2} \tag{58}$$

Using this equation we now have tighter bounds for the error. To make easier use of the bounds

$$\text{Let } \epsilon''(n) = \frac{1}{24n^2} \tag{59}$$

$$0 < \epsilon''(n+1) < H_n - \ln(n + \frac{1}{2}) - \gamma < \epsilon''(n) \tag{60}$$

Substituting into Eq. 38 we have

$$\ln(n+k+\frac{1}{2}) + \gamma - (\ln(n+\frac{1}{2}) + \gamma + \epsilon''(n)) < \frac{c}{v} < \ln(n+k+\frac{1}{2}) + \gamma + \epsilon''(n+k) - (\ln(n+\frac{1}{2}) + \gamma) \tag{61}$$

Clearly, $\epsilon''(n+k) < \epsilon''(n)$.

$$\ln(n+k+\frac{1}{2}) - \ln(n+\frac{1}{2}) - \epsilon''(n) < \frac{c}{v} < \ln(n+k+\frac{1}{2}) + \epsilon''(n) - \ln(n+\frac{1}{2}) \tag{62}$$

$$\frac{n+k+\frac{1}{2}}{n+\frac{1}{2}} e^{-\epsilon''(n)} < e^{\frac{c}{v}} < \frac{n+k+\frac{1}{2}}{n+\frac{1}{2}} e^{\epsilon''(n)} \tag{63}$$

$$(n+k+\frac{1}{2})e^{-\epsilon''(n)} < (n+\frac{1}{2})e^{\frac{c}{v}} < (n+k+\frac{1}{2})e^{\epsilon''(n)} \tag{64}$$

Solving each inequality separately, we have

$$(n+k+\frac{1}{2})e^{-\epsilon''(n)} < (n+\frac{1}{2})e^{c/v} \qquad (n+\frac{1}{2})e^{c/v} < (n+k+\frac{1}{2})e^{\epsilon''(n)} \tag{65}$$

$$n+k+\frac{1}{2} < (n+\frac{1}{2})e^{c/v}e^{\epsilon''(n)} \qquad n+k+\frac{1}{2} > (n+\frac{1}{2})e^{c/v}e^{-\epsilon''(n)} \tag{66}$$

$$k < (n+\frac{1}{2})(e^{c/v}e^{\epsilon''(n)} - 1) \qquad k > (n+\frac{1}{2})(e^{c/v}e^{-\epsilon''(n)} - 1) \tag{67}$$

A better approximation for k than Eq. 53 is

$$k \approx (n+\frac{1}{2})(e^{c/v} - 1) \tag{68}$$

The error range now is

$$(n+\frac{1}{2})(e^{c/v}e^{\epsilon''(n)} - 1) - (n+\frac{1}{2})(e^{c/v}e^{-\epsilon''(n)} - 1) = e^{c/v}(n+\frac{1}{2})(e^{\epsilon''(n)} - e^{-\epsilon''(n)}) \tag{69}$$

While our previous approximation gave constant bounds on the error range as n grew. This approximation can be shown to converge to a single value as n grows.

$$\text{Consider, } \lim_{n\to\infty}(n+\frac{1}{2})(e^{\epsilon''(n)} - e^{-\epsilon''(n)}) = \lim_{n\to\infty}(n+\frac{1}{2})(e^{\frac{1}{24n^2}} - e^{-\frac{1}{24n^2}}) \tag{70}$$

$$\text{Substitute } x = \frac{1}{24n^2} \tag{71}$$

$$\lim_{x\to 0} \frac{e^x - e^{-x}}{\sqrt{24x}} - \frac{e^x - e^{-x}}{2} = 0 - 0 = 0 \tag{72}$$

The second term goes to 0, and, using L'Hopital's rule, the first term also goes to 0. So for sufficiently large n, the error range converges to 0.

What happens for small n? With our previous approximation, the error range for $n = 1$ was approximately $1.04e^{c/v}$, which had an upper bound of approximately 2.8, when $c = v$. With the improved approximation the error range at $n = 1$ is less than $0.13e^{c/v}$, which in the worst case, is less than 0.34. Calculating k to the nearest integer will be correct with very high probabilty.

Appendix B. Proof of Unique Global Maximum for Segregated Schedule Utility

From Equation 15 we have the following equation for the utility of a segregated schedule of length Z with x cooperations followed by $Z - x$ defections. We may express the total utility of such a schedule as

$$U_{seg}(Z, x) = U = (v - c)x + \sum_{i=x}^{Z-1} v\frac{x}{i} \tag{73}$$

To prove there can only be one unique value of x that maximizes $U_{seg}(Z, x)$ for a given Z, we will take the second derivative with respect to x and show that it only takes on negative values for $0 \leq x \leq Z$.

We begin by simplifying Eq. 73.

$$U = (v - c)x + \sum_{i=x}^{Z-1} v\frac{x}{i} \tag{74}$$

$$= (v - c)x + vx\Big(\sum_{i=1}^{Z-1} \frac{1}{i} - \sum_{i=1}^{x-1} \frac{1}{i}\Big) \tag{75}$$

We now have the difference of two finite harmonic sums. A finite harmonic sum can be expressed analytically as

$$H_n = \gamma + \psi_0(n + 1) \tag{76}$$

where γ is the Euler-Mascheroni constant and $\psi_0(n + 1)$ is the digamma function [24]. Substituting in for the series gives us

$$U = (v - c)x + vx(\gamma + \psi_0(Z) - (\gamma + \psi_0(x))) \tag{77}$$

Next, we simplify and take two derivatives. The derivative of $\psi_0(z)$ is $\psi_1(z)$ and similarly the derivative of $\psi_1(z)$ is $\psi_2(z)$, where $\psi_1(z)$ and $\psi_2(z)$ are polygamma functions [25].

$$U = (v - c)x + vx(\cancel{\gamma} + \psi_0(Z) - \cancel{\gamma} - \psi_0(x)) \tag{78}$$

$$= (v - c)x + vx(\psi_0(Z) - \psi_0(x)) \tag{79}$$

$$= (v - c)x + v\psi_0(Z)x - vx\psi_0(x) \tag{80}$$

$$\frac{dU}{dx} = (v-c) + v\psi_0(Z) - v\psi_0(x) - vx\psi_1(x) \tag{81}$$

$$\frac{d^2U}{dx^2} = -v\psi_1(x) - v\psi_1(x) - vx\psi_2(x) \tag{82}$$

$$= -v(2\psi_1(x) + x\psi_2(x)) \tag{83}$$

(84)

A polygamma function $\psi_n(z)$ can be written as follows [25]:

$$\psi_n(z) = (-1)^{n+1} n! \sum_{k=0}^{\infty} \frac{1}{(z+k)^{n+1}} \tag{85}$$

Applying Eq. 85 to Eq. 84 and simplifying:

$$\frac{d^2U}{dx^2} = -v(2\psi_1(x) + x\psi_2(x)) \tag{86}$$

$$= -v\left[2\left(\sum_{k=0}^{\infty} \frac{1}{(x+k)^2}\right) + x\left(-2\sum_{k=0}^{\infty} \frac{1}{(x+k)^3}\right)\right] \tag{87}$$

$$= -2v\sum_{k=0}^{\infty}\left(\frac{x+k}{(x+k)^3} - \frac{x}{(x+k)^3}\right) \tag{88}$$

$$= -2v\sum_{k=0}^{\infty}\left(\frac{k}{(x+k)^3}\right) \tag{89}$$

Notice that for any valid value of x, $0 < x \leq Z$, the summation is purely positive. Therefore, the second derivative of $U_{seg}(Z, x)$ must be negative in that same range.

Appendix C. Estimating Optimal Schedule for Fixed Number of Transactions

From Equation 9 we know that, given a number of completed transactions n and a cost/valuation ratio c/v, we can calculate how many additional transactions k are needed so that the utility from cooperating or defecting on the $n+1$ turn is equal. Consequently, a seller benefits more from defecting if she participates in less than k additional transactions, and benefits more from cooperating if she participates in more than k additional transactions. Therefore, for a given number of total transactions Z, we can determine how many cooperations are optimal by computing n and $k+1$ respectively from $Z = n + 1 + k$, substituting in Equation 9 for k, then solving for n.

$$Z = n + 1 + k \tag{90}$$

$$Z = n + 1 + (n + \frac{1}{2})(e^{c/v} - 1) \tag{91}$$

$$Z = \cancel{n} + 1 + ne^{c/v} - \cancel{n} + \frac{1}{2}e^{c/v} - \frac{1}{2} \tag{92}$$

$$n = (Z - \frac{1}{2})e^{-c/v} - \frac{1}{2} \tag{93}$$

We now have the optimal number of cooperations in terms of Z, the total number of transactions, $n_C(Z)$. Now subtracting the value of n in Eq. 93 from Z gives us the optimal number of defections in terms of Z, $n_D(Z)$.

$$n_D(Z) = Z - ((Z - \frac{1}{2})e^{-c/v} - \frac{1}{2}) \tag{94}$$

$$= Z - (Z - \frac{1}{2})e^{-c/v} + \frac{1}{2} \tag{95}$$

$$= (Z - \frac{1}{2}) + \frac{1}{2} - (Z - \frac{1}{2})e^{-c/v} + \frac{1}{2} \tag{96}$$

$$= (Z - \frac{1}{2})(1 - e^{-c/v}) + 1 \tag{97}$$

The previous equations allow for real numbered values. Because we are interested in integer values we must apply proper integer conversions. For a fixed number of transactions Z, the utility optimal number of cooperations and defections, respectively, are

$$n_C(Z) = \left\lceil (Z - \frac{1}{2})e^{-c/v} - \frac{1}{2} \right\rceil \tag{98}$$

$$n_D(Z) = \left\lfloor (Z - \frac{1}{2})(1 - e^{-c/v}) + 1 \right\rfloor \tag{99}$$