On Experimental Equilibria Strategies for Selecting Sellers and Satisfying Buyers*

Claudia V. Goldman†
Dept. of Computer Science
University of Massachusetts
Amherst, MA 01002
clag@cs.umass.edu

Sarit Kraus
Dept. of Math and CS,
Bar Ilan University,
Ramat Gan, 52900 Israel
sarit@cs.biu.ac.il
and
Inst. Adv. Comp. Studies,
University of Maryland,
College Park, MD 20742 USA

Onn Shehory
IBM Research Labs in Israel,
c/o Tel Aviv site
Haifa University,
Mount Carmel, Haifa
31905 Israel
onn@il.ibm.com
Tel: 972-3-640-1682
Fax: 972-3-640-1631

Abstract

We consider marketplaces where buyers and sellers iteratively encounter to trade. Given some specific trade conditions, the question that we address is what strategies should buyers and sellers use to maximize gains. We focus on electronic markets where supply shortages are common. Under such market conditions sellers can only satisfy a subset of the purchase-orders they receive from buyers. Consequently, some buyers may become discontented and they may be motivated to migrate to other sellers in proceeding encounters.

Beneficial purchase-order selection as well as seller selection require, respectively, seller and buyer strategies. Analytical computation of stable profiles of such strategies is infeasible in the enviroments we examine. We hence devise a new methodology for studying strategic equilibria. We introduce specific equilibria strategy profiles to be implemented by automated trade agents. The main conclusions of our study are that automated sellers will benefit most by randomly selecting the purchase-orders of their buyers to be satisfied. Additionally, such sellers will not benefit from learning the buyers' typical order size. Moreover, automated buyers will maximize their benefits by re-issuing purchase-orders with sellers that satisfied them, fully or partially, in the past.

Keywords: multi-agent market; agents' strategies; experimental equilibrium.

*This research was supported in part by NSF under grant No.IIS9820657 and IIS0208608 and by IBM Research. Preliminary results were reported in CIA-01 and MAAMAW-01. Correspondence author: Onn Shehory

†This work was performed while the first author was a post-doctoral fellow at Bar Ilan University.

1 Introduction

Marketplaces, either traditional or electronic, are sites in which buyers and sellers meet in order to trade. In monetary-based trade, goods held by sellers are sold in exchange for monetary funds held by buyers. Although the task of a seller is to sell its merchandise and the task of a buyer is to get the goods it needs, doing so, both sellers' and buyers' main goal is to maximize their gains. Seeking such maximization, buyers and sellers may practice strategic behavior. This imposes a practical question: given some specific trade conditions, what strategies should buyers and sellers use to maximize gains? Game theory - the discipline in which strategic behavior is studied analytically - provides very few practical solutions to this question. This is mainly because the problem is too complex to solve analytically and the strategy space is in many cases intractable. Such is the case in several electronic types of trade. Prominent examples can be found both in B2C (business-to-consumer) and B2B (business-to-business) marketplaces. For instance, in a competitive B2C commodity market (e.g., books), where prices of goods are virtually the same across seller sites, and sellers hold limited stocks to reduce their storage and financing expenses, stock shortages result in buyers not being serviced. A similar phenomenon is present in B2B MRO (Maintenance, Repair, Operation) markets (e.g., office supplies), where prices are competitive and sellers' profits are lean.

Under the trade conditions referred to above, where sellers may, in times, be able to address only part of the orders placed by buyers at their site, the sellers are in need for strategies that will minimize the damage from the orders they received but eventually rejected. Buyers can benefit as well from strategies that increase their chances of being serviced. The goal of this research is to find such strategies. In particular, we seek strategies that, given the trade conditions referred to above, maximize the gains of both sellers and buyers, and maintain a stable market. As stated above, computing these strategies analytically is infeasible. Therefore, we have developed an experimental methodology as well as a simulation tool that implements it. We use these methodology and tool to find the sought strategies.

The strategies we seek are different from strategies applicable to traditional trade. The reason for this is that, in difference from traditional trade, electronic markets introduce a combination of conditions that does not exist in other markets. This combination raises challenging problems which are unique to electronic trade. The electronic interaction among buyers and sellers in the markets we study includes the following characteristics:

1. Buyers can keep their anonymity (e.g., via anonymizers and third party brokers).
2. Buyers do not know about other buyers that participate in the trade, even when those visit the same store concurrently.

3. Buyers do not know the actual true stock level of the seller.
4. Sellers are uncertain regarding the number of buyers that will approach their site and the volume of their orders. This problem is more intense in electronic markets than it is in traditional trade. It results from liquidity and dynamism of the activity on the web, where sites are easily accessible by buyers from diverse geographical locations, and the cost of switching from one site to another is very low.
5. The pace of trade is high. This results from the automation of seller sites as well as their distribution. The pace shall be further accelerated when the buyer's side will be automated via buyer agents.

The above characteristics result in several conditions. For instance, attribute 1 results in the sellers' inability to identify specific anonymous buyers. Thus, sellers may be unable to benefit from identity-based long-term customer relation management. Attribute 3 results in the sellers' ability to manage stocks and orders at their own discretion to maximize benefits. Attributes 4 and 5 above result in dynamic variations in demands at sellers' sites, which impose occasional shortages in stock.

Instantaneous supply shortages are rather common in B2C and B2B MRO electronic commerce. One reason for this, which is unique to electronic trade, is that inventory verification for a given purchase-order is performed only after the order is complete. Consequently, concurrent preparation of purchase-orders by multiple buyers may result in a seller being unable to satisfy the orders. This is in spite of the seller (implicitly or explicitly) declaring that the requested products are available for sale.

In such cases of instantaneous supply shortages, sellers must decide which orders to satisfy and which ones not to. This will leave some buyers unsatisfied. Buyers unsatisfied with one seller may be motivated to approach another seller in proceeding encounters. This behavior is referred to as punishment (for not providing the expected supply).

We have studied electronic markets where such trade conditions hold. The major contribution of this study is twofold:

- We introduce a new methodology for studying strategic equilibria in environments where analytical computation of these is infeasible. In particular, this methodology is applicable for B2C and B2B MRO electronic markets.

- We find and present specific equilibria strategy profiles to be implemented by automated trade agents. These agent-oriented strategies can be used by humans who trade in electronic markets as well.

Based on the study of strategy profiles we performed, the main conclusions are the following. Automated sellers will benefit most by randomly selecting the purchase-orders of their buyers in order to satisfy these orders. Additionally, such sellers will not benefit from learning the buyers' typical order size. Further, automated buyers will maximize their benefits by re-issuing purchase-orders with sellers that satisfied them, fully or partially, in the past.

The rest of the paper is organized as follows. In Section 2 we present other studies relevant to our research. Then, in Section 3, we describe the model of the electronic market we examine and its specific settings. We proceed with presenting the strategies examined in this market in Section 4. Following, in Section 5, we provide details regarding the simulation methodology and implementation used for our experiments. The results of these experiments, in which we study strategic interaction between electronic sellers and buyers, are presented in Section 6. Section 7 closes with conclusions and discussion of future work.

2 Related Work

Work in the field of computational scheduling (e.g., [9]) addresses the distribution of a set of tasks to a set of computers. The problem handled in this paper may be viewed as a distributed scheduling problem. A number of buyers needs to be distributed among a number of sellers such that the sellers remain with the minimal number of unsold items in their stocks, and the buyers are satisfied as much as possible. A prominent difference between scheduling solutions and electronic markets' solutions is that the former attempt to minimize some global objective functions, whereas in the latter, there is no such global function. In electronic markets buyers are represented by intelligent agents whose purpose is to each maximize its own gain and they decide rationally based on their strategy of action. Sellers are self-interested and do not act under a central control either. Another difference is that, unlike scheduling solutions, our analysis needs to take into account the buyers' reaction to the sellers' actions. As a result of these differences between scheduling problems and electronic markets problems, we need to provide other solutions for the latter.

Research in game theory and economics has addressed issues of competition among buyers and among sellers that need to choose which buyer orders to satisfy. To our best knowledge, none of these studies addressed agent strategic behavior in large markets. For instance, Vincent [11] studies *a single seller* and the way in which the seller selects one buyer from among a group of buyers, where all of the buyers are interested in buying the *single good* the seller wishes to sell. We analyze a market where multiple sellers interact repeatedly with multiple buyers. Each buyer places its purchase-orders sequentially. At each encounter, some sellers cannot fulfill some of the buyers' purchase-orders on time.

A market in which sellers face an infinite number of anonymous and typed buyers is studied in [1]. There, buyers compete since only one of them will be satisfied by the seller. The question answered there was what price the winning buyer will pay. We consider a different setting, where multiple products are sold, and we address a different question, as our goal is to identify, by examining various strategy profiles, those strategies that agents should follow to maximize their gains, assuming that prices are given.

In market settings similar to those we study, when the buyers need to choose which sellers to approach, they may consider the sellers' reputations. Reputation was studied as a function (e.g., [13]), as a social mechanism for avoiding untrustworthy parties [12], and as an expectation for quality [10]. Reputation is usually built over time and its use is appropriate for long-term market interactions. In our case, interaction are mid-term to short-term. Hence, the use of reputation systems is impractical. Yet, the buyers' strategy may change according to the service received from the sellers. This has some similarity to the use of reputation, excluding the aggregation effect.

A market where buyers evolve into customers is studied in [2]. In particular, a model of the buyers is analyzed, given that the sellers have a fixed capacity and receive a fixed and equal demand from all buyers. There, it is assumed that each seller, subject to its capacity, serves as many buyers as it can to entirety, and partially satisfies one buyer. It is also assumed that the game lasts for a number of periods greater than, or equal to, the number of sellers in the market. It is shown that conditional loyalty leads to an equilibrium. That is, buyers will return to sellers that have served them in the last trading interaction, and sellers serve the buyers they have served in the past. That study seems similar to ours, as it addresses multiple sellers and multiple buyers. Yet, in that research the sellers' strategy was set, and only the buyers reaction to this preset strategy was examined. In our research, both buyers and sellers can follow one out of a set of strategies, and we investigate equilibria among these strategies, without imposing a specific strategy on either buyers or sellers.

A critical approach to game theory applied to Internet games appears in [3]. The authors claim that these games cannot be analyzed in the framework of Nash equilibrium because there is lack of information (the buyers do not know who the other buyers are and what they do), the players do not apply optimal procedures, there is no synchronization, and the mid-term is the one that is relevant. The authors suggest *reasonable learners* algorithms, whose important features are optimality (i.e., the algorithms should enable the agent to get the optimal gain), monotonicity (if the payoff of performing an action increases, the probability of choosing this action cannot decrease) and responsiveness. In our work, we evaluate actions based on the notion of experimental equilibrium, which enables us to find strategy profiles that maximize the expected utility of the agents given the behavior of the other market participants.

Kephart et al. [7, 8] have also taken the game theoretic approach to analyze electronic markets. Their focus, though, is on dynamic price changes in electronic markets. In this paper, we assume that the time range is short and therefore prices remain static. Our focus is on the interactions between the agents' strategies.

3 The Simulation Model

A general analytical solution to the strategic agent interaction problem that we address does not seem amenable. To provide a solution nevertheless, we opted for an experimental approach. Hence, we have developed a simulation model—Simulation of Electronic Market Interactions (SEMI)—which was also implemented and utilized to study possible strategies for the sellers and the buyers in the markets of interest. Running simulations with the SEMI system enables us to empirically test the effects of various settings on the utilities obtained by the electronic buyers and sellers. Experiments with SEMI can (and do) unravel experimental equilibria between buyers' and sellers' strategies.

The SEMI system can be used for a broad range of settings, however we use it to model and simulate the specific problem of interest. That is, we simulate sellers that hold limited stocks and buyers that each places a purchase-order with one seller at each encounter. Without loss of generality, we refer to a single type of good[1] to be sold by all of the sellers. We consider a market of repeated encounters between finite sets of buyers and sellers. Buyers approach sellers and submit requests for buying goods. At each encounter, a seller may receive requests that cumulatively exceed the quantity of the good available in its stock. It thus needs to decide which requests to satisfy and which ones to leave unsatisfied.

In our model, there is a set $\mathcal{B}$ of buyers and a set $\mathcal{S}$ of sellers. A SEMI simulation consists of a sequence of K encounters between the sellers and the buyers. In each encounter, k, a seller $S^j \in \mathcal{S}$ holds in its stock a quantity ST^{jk} of the good for sale. Each buyer $B^i \in \mathcal{B}$ is associated with a type TY^i which indicates the average size of its purchase-order. We denote B^i's purchase-order at encounter k by PO^{ik} or, when this order is intended for a specific seller S^j, by PO^{ik}_j. The purchase-order, PO^{ik}, specifies the number of units of the good that B^i would like to buy at encounter k. To simplify the analysis, we assume that this number is an integer.[2] At a given encounter k, B^i submits its purchase-order PO^{ik}_j to some seller, S^j. S^j can sell B^i all, part or none of the quantity specified in its purchase-

[1]Goods of various types can be handled in the same way as long as there are no interdependencies among them.

[2]This assumption is not restrictive, as the majority of the B2C and B2B MRO goods are sold in whole units.

order. B^i's utility from this transaction is proportional to the portion of the purchase-order supplied to it. In our model, a seller gains from the sale of each unit and it incurs a cost for each unsold unit held in its stock at the end of each encounter. Thus, a seller's utility at a given encounter is proportional to the quantity of the good it sells at that encounter, and is inversely proportional to the quantity of the good left in its stock at the end of the encounter. A seller would incur additional costs for revealing the buyers' types. The utility functions of the agents explicitly express these costs. The details of the utility functions and their formal expression are provided later in this section.

Our model includes several assumptions that commonly hold in B2C and B2B MRO electronic markets or in a subset thereof. We assume that all of the sellers sell the same type of a good, where quality, price and units are the same as well. The conditions of this assumption hold in a subset of the electronic markets. They typically hold in competitive commodity markets such as, e.g., books, appliances, office supplies and others, where the ease of price comparison and the competition among sellers result in relatively uniform prices and quality. Our model refers to sellers and buyers that transact repeatedly in fixed-price, catalog-based, electronic stores, where prices remain static within short to intermediate periods of time. Our solution concentrates on such periods of time, hence we can safely assume static prices.[3] To simplify our analysis we assume that, at each encounter, each buyer is associated with only one purchase-order, and this purchase-order is valid only for that encounter. Such a condition holds only in some electronic markets, however it is rather simple to relax this assumption. For instance, to model purchase-orders that persist along multiple encounters, all one needs to do is to allow re-issuing of expired purchase orders in the encounter proceeding their expiration.

Following the settings and assumptions above, we present the details of buyers' and sellers' utility functions as well as other controllable parameters of our simulation model.

- **The utility function of the sellers** — We examine risk-prone, risk-neutral and risk-averse sellers. The utility function varies according to the risk attitude. Given a risk-neutral utility function U, the risk-prone utility function is convex w.r.t. U, and the risk-averse utility function is concave w.r.t. U [4].

  The utility function of a risk-neutral seller S^j in a given encounter k, when it receives r purchase-orders, is

$$U^j(PO_j^{1k}, \ldots, PO_j^{rk}) = G \cdot \sum_{i=1}^{r} sat(PO_j^{ik}) - C_s \cdot (ST^{jk} - \sum_{i=1}^{r} sat(PO_j^{ik})) - C_t \cdot r^* \quad (1)$$

[3]See reference to work by Kephart as explained in Section 2 for dynamic pricing strategies.

G is the gain for one sold unit; the function $sat(PO_j^{ik})$ returns the size of satisfied portion of PO_j^{ik}; C_s is the cost of holding one unsold unit for one encounter; C_t is the cost for obtaining the type of one buyer; r^* is the number of times the seller bought such type information in the given encounter.[4]

The utility function of a risk-prone seller is the following convex function:[5]

$$U^j_{convex}(PO_j^{1k}, \ldots, PO_j^{rk}) = exp(\frac{G \cdot \Sigma_{i=1}^r sat(PO_j^{ik}) - C_s * (ST^{jk} - \Sigma_{i=1}^r sat(PO_j^{ik}))}{ST^{jk}}) \tag{2}$$

The utility function of a risk-averse seller is the following concave function:

$$\begin{aligned} U^j_{concave}(PO_j^{1k}, \ldots, PO_j^{rk}) &= exp(\frac{C_s * ST^{jk}}{Av}) - \\ & exp(\frac{G\Sigma_{i=1}^r sat(PO_j^{ik}) - C_s * (ST^{jk} - \Sigma_{i=1}^r sat(PO_j^{ik}))}{-Av}) \end{aligned} \tag{3}$$

where Av a positive number that affects the level of averseness of the sellers such that smaller Av refers to higher averseness and vice versa.

- **The utility function of the buyers** — We distinguish between conceding buyers and non-conceding buyers. A conceding buyer's utility is proportional to the part of its purchase-order that was fulfilled. The utility function is given by $U^i(PO^{ik}) = sat(PO^{ik})/PO^{ik}$. Non-conceding buyers are buyers whose willingness to accept partial purchase-order fulfillment is lesser than the willingness of other (conceding) buyers to do so. This is expressed in their utility function, where partial satisfaction of purchase-orders intensifies the loss of gains. A non-conceding buyer B^i's utility is given by the function
$$U^i_{nc}(PO^{ik}) = (sat(PO_j^{ik}) - (PO_j^{ik} - sat(PO_j^{ik})) * LS)/PO_j^{ik} \tag{4}$$
where LS is the factor that reflects the level of buyer dissatisfaction with partial purchase-order fulfillment. Note that here, the utility of a partially satisfied buyer

[4]Note that r^* is only relevant to the strategies that consider buyers' type. In our experiments, these strategies were found to provide poor gains compared to other strategies. Therefore, the use of these strategies is of little importance and r^* is omitted from the following utility functions.

[5]It is based on the non-normalized utility function of the risk-neutral sellers given by $G\Sigma_{i=1}^r sat(PO_j^{ik}) - C_s * (ST^{jk} - \Sigma_{i=1}^r sat(PO_j^{ik}))$.

Notation	Description
k	The k^{th} encounter.
K	The number of encounters.
$\mathcal{B}=\{B^i\}, 1 \leq i < \infty$	The set of buyers.
$\mathcal{S}=\{S^j\}, 1 \leq j < \infty$	The set of sellers.
ST^{jk}	The stock that seller S^j holds at the beginning of encounter k.
TY^i	The average purchase-order size of buyer B^i.
PO_j^{ik}	The purchase-order placed by buyer B^i to seller S^j at encounter k.
$sat(PO_j^{ik})$	The actual deal satisfied by seller S^j after buyer B^i requested PO_j^{ik}.
C_s	The cost of holding one unsold item in stock for one iteration.
C_t	The price that a seller needs to pay to reveal one buyer's type.
G	The gain of a seller from selling one unit of the good.

Table 1: Summary of the SEMI's Notations

may be nullified.[6] As the value of LS increases, the buyer's level of non-concession increases too.

- **Stock sizes** — We distinguish between homogeneous and heterogeneous markets. In homogeneous markets, all sellers hold equal-sized stocks at the beginning of each encounter (i.e., for any $k, k' \in I\!N$, $ST^{jk} = ST^{jk'}$). In heterogeneous markets, the sellers may hold stocks of different sizes.

- **The size of the market** — Smaller markets may behave differently from larger ones, in the sense that sellers and buyers may need to use different strategies to maximize their gains. Our model allows us to populate the market with buyers' and sellers' populations of various sizes. Utilizing this flexibility, we have examined markets of several sizes. By this, we were able to study similarities and differences in strategic behavior as affected by the size of the market.

The utility functions and the parameters presented above are used in our experiments. The notations of this section are summarized in Table 1. These will be used in following sections to describe strategies (Section 4.1) and experiments (Section 6).

4 Sellers' and Buyers' Strategies

The model presented above provides the framework in which the trade interactions of interest should take place. Given this framework, agents would interact strategically to achieve their trade objectives. Since we use a simulated agent system for our study, we need to provide the

[6] For example, when $LS = 0.5$, a buyer for which less than a third of the purchase-order was satisfied will gain zero, which is, utility-wise, equal to not being satisfied at all.

strategies to be used by the agents within the simulator. We denote by Σ^s the set of strategies to be used by the sellers, and by Σ^b the set of buyers' strategies. It is preferable that the strategies we provide represent the strategy space, however since this space is intractable, we have concentrated on a set of representative strategies. We later show that the selected strategies are beneficial and enhance stability.

As stated in Section 1, the markets we study are characterized by sellers that, as a result of stock shortages, may be unable to satisfy some of the requests submitted to them. Therefore, we concentrate on buyer strategies that allow punishment behaviors, and seller strategies that consider a variety of decision policies for selecting buyers to be satisfied.

4.1 Sellers' strategies

In our model, given a set of purchase-orders received by a seller, a seller's strategy specifies which portion of each purchase-order to supply. The strategy could be affected by the following issues: (i) the arrival time of the purchase-order at the seller's site (within an encounter); (ii) the size of the purchase-order; and (iii) the type of the buyer that submitted the purchase-order, when it is possible for the sellers to obtain this information.

We developed representative strategies to test which equilibria will result from encounters between buyers and sellers. We categorize the strategies according to the way in which they use the available information. For simplicity, a strategy does not use historical information. This is reasonable in situations where the buyers are unrecognizable. Note that, when the requests to a seller, cumulatively, do not exceed its stock, the differences between the strategies become unimportant, as the seller agent will satisfy all the requests, regardless of the use of a specific strategy. Only when the requests cumulatively exceed the seller's stock, the details of the strategy indeed affect the behavior of the agent that implements it. Below are the details of the strategies. In the naming of these strategies, we add the prefix letters O, D and R, to indicate that the strategy implements, respectively, size ordering, proportional distribution and type recognition.

Below we present the strategies, partitioned into categories:
The **uninformed seller** category includes strategies in which the seller does not consider the size of the purchase-orders (issue (ii) above) nor the buyers' type (issue (iii) above). In this category we consider a simple seller strategy as follow.

1. Random Seller (in short RandS)— According to this strategy, a seller randomly chooses the purchase-orders to be fulfilled from those that were submitted to it. The seller attempts to completely fulfill all of the requests, however as a result of its stock being limited in size it may end up not fulfilling some of the requests. The seller sequentially

selects, by random, purchase-orders. Each of these is fully satisfied (except for, possibly, one purchase-order which may be only partially satisfied), until its stock is exhausted or there are no more requests to fulfill. The RandS strategy is equivalent to a First-In-First-Served (FIFS) strategy in cases where the order of arrival of buyers' requests at sellers' sites does not depend on characterizing traits of the buyers (e.g., their typical order size).

The **greedy seller** category includes strategies that are affected by the size of the purchase-orders (issue (ii) above).

1. Ordered Purchase-Orders (in short OPO) — According to this strategy, the fulfillment of the buyers' purchase-orders is performed in a decreasing order of their sizes.
2. Distributed Purchase-Orders (in short DPO) — According to this strategy, a seller (partially) satisfies the purchase-orders proportionally to their size. That is, if S^j receives r purchase-orders in encounter k, $PO_j^{1k}, \ldots, PO_j^{rk}$, then each buyer B^i will be supplied with $\lfloor \frac{PO_j^{ik}}{\sum_{l=1}^{r} PO^{lk}} \cdot ST^{jk} \rfloor$. The remainder of this distribution, i.e., $ST^{jk} - (\sum_{h=1}^{r} \lfloor \frac{PO_j^{hk}}{\sum_{l=1}^{r} PO^{lk}} \cdot ST^{jk} \rfloor)$, is allocated to one buyer, selected randomly.

The **intelligent seller** category includes strategies in which the seller uses the buyers' type (issue (iii) above) for deciding upon purchase-orders to be fulfilled.

1. Ordered Types (OType) and Ordered Recognized Types (ORType) — According to these strategies, the fulfillment of the buyers' purchase-orders is performed in a decreasing order of the type of the buyers. R in the prefix denotes recognizable buyers (referring to type recognition and not identity recognition). In our model, sellers that are interested in the information regarding the type of a buyer need to pay for it. When buyers are recognizable (i.e., their identity is not hidden), a seller needs to pay only once for this information. In such a case, the ORType strategy is relevant. When buyers are not recognizable, the a seller needs to pay for this information in each encounter it would like to use it. In such a case, the OType strategy is relevant.
2. Distributed Types (DType) and Distributed Recognized Types (DRType) — According to this strategy, a seller satisfies the purchase-orders proportionally to the buyers' types. That is, if S^j receives r purchase-orders in encounter k, $PO_j^{1k}, \ldots, PO_j^{rk}$, then for each

buyer B^i, the seller computes $po_i = \lfloor \frac{TY^i}{\sum_{l=1}^{r} TY^l} \cdot ST^{jk} \rfloor$. If $po_i \geq PO_j^{ik}$ then po_i is set to PO_j^{ik}. The remainder of the seller's stock is allocated to buyers in the same way as in the DPO strategy. The inclusion of R in the prefix (DRType) refers to the case of recognizable buyers. Payment for type are applicable following the same guidelines as in the ORType and OType strategies.

The sellers' behavior when following the strategies presented above is demonstrated by examples in Appendix A.

4.2 Buyers' strategies

Under the assumptions made in this paper, a buyer has very limited information (relevant to the trade) on the sellers. Note that limited information as we assume is common in real B2C and MRO electronic markets. At each encounter, the buyer knows what portion of its submitted purchase-order was satisfied. Given such a history of past encounters, a buyer should decide which seller to approach in the current encounter. In this paper, for simplicity, we focus on strategies that take into consideration only the history of the last encounter.[7] The list of buyer strategies follows.

1. Random Buyer (in short RandB) — According to this strategy, the buyer randomly selects a seller for submitting a request.
2. Loyal — According to this strategy, a buyer B^i will first check whether the seller S^j with which it has placed a purchase-order at encounter k has *completely* or at least *partially* satisfied it at that encounter. If S^j has (partially) satisfied it, at encounter k, B^i returns to seller S^j at encounter $k+1$. Otherwise, at encounter $k+1$, B^i randomly selects a seller from $\mathcal{S}$ to which it submits its purchase-order (this random selection may be S^j as well).
3. Loyal and Punish (in short LoyalP) — According to this strategy (in similarity to the Loyal strategy), a buyer B^i will first check whether the seller S^j with which it has placed a purchase-order at encounter k has *completely* or at least *partially* satisfied it at that encounter. If S^j has (partially) satisfied it at encounter k, B^i returns to seller S^j at encounter $k+1$. In difference from the Loyal strategy, if B^i's order was not satisfied at all, at encounter $k+1$ B^i randomly selects a seller from $\mathcal{S} \setminus \{S^j\}$ (thus not approaching S^j at encounter $k+1$).

[7]Long histories are more commonly taken into account in long-term interactions between buyers and sellers. In these cases, contracts are usually signed when the supply and the demand are known in advance. This is not the case we study in this research, in which we focus on short-term interactions.

4. Loyal Weak (in short LoyalW) — According to this strategy, a buyer B^i will first check whether the seller S^j with which it has placed a purchase-order at encounter k has *completely* satisfied it at that encounter. If S^j has fully satisfied it at encounter k, B^i returns to seller S^j at encounter $k+1$. Otherwise, at encounter $k+1$, B^i randomly selects a seller from $\mathcal{S} \setminus \{S^j\}$ (thus not approaching S^j at encounter $k+1$).
5. Probabilistic Buyer (in short Prob) — According to this strategy, a buyer B^i that has approached seller S^j at encounter k with a purchase order PO_j^{ik}, will approach S^j at encounter $k+1$, with a probability of $sat(PO_j^{ik})/PO_j^{ik}$. This probability—the ratio between the portion of the purchase-order satisfied and the whole purchase-order—expresses the level of satisfaction of a buyer. B^i will approach each of the other sellers with a probability of $(1 - sat(PO_j^{ik})/PO_j^{ik})/(|\mathcal{S}| - 1)$.

When following any of the strategies above, except for the RandB strategy, a buyer will return to a seller that has completely satisfied it. Yet, when partially satisfied or not satisfied at all, a buyer may *punish* the seller it has approached at encounter k, by not returning to this seller at encounter $k+1$. The severity of the punishment is expressed in the following order of the strategies: on the one extreme, LoyalW is the most punishing strategy because even if a seller has partially satisfied a buyer, this buyer will not return to the seller at the next encounter. The probabilistic strategy is less punishing than LoyalW since there is still a positive probability for a buyer to return to a partially satisfying seller. LoyalP is more punishing than Loyal and less punishing than LoyalW. RandB induces a buyer to choose a seller in a random way, with no regard to the seller's behavior. We demonstrate the behavior of the buyers when they follow each of the aforementioned strategies in Appendix B. Our hypothesis was that not returning to a seller that has not fully satisfied a buyer, would be the best strategy for the buyer. We found out in the experiments that this is not always the case.

5 The Methodology of the Simulation-based Solution

The utility of an agent trading in an electronic market is influenced by the other agents' actions in that market. The analysis of such influences usually resides in the field of game theory, hence our methodology relies on game-theoretic concepts as well. In particular, we study strategy profiles and their stability. However, in difference from the classical game-theoretic approach, we perform our study via simulations. In the simulations, the expected utility of buyers and sellers is computed subject to various market settings and the use of different strategies. Since in markets of the type studied here agents are self-interested and competitive, there is no single evaluation criterion that fully captures the preferences of all market participants. Some solutions might be optimal for some of the agents, but disadvantageous for others. Our study seeks solutions, in which each agent maximizes its utility *given* the behavior of the other agents. Solutions of this type provide strategy profiles that inherently give the agents no incentive to deviate from them.

Given the two sets of strategies Σ^s and Σ^b as described earlier, and with such a game-theoretic approach taken, we seek strategy profiles that are in equilibrium. A strategy profile F is a set of strategies, one for each buyer and one for each seller from the relevant sets of strategies. To find strategy profiles which are in equilibrium, we compute the average utility of each agent in the market, utilizing the SEMI simulator, and use this computed utility as an estimation of the agent's expected utility. Once the utilities are available, we examine strategy profiles in conjunctions with the utilities they yield, seeking profiles that are in *experimental equilibrium*.

Definition 1 (Experimental Equilibrium) *A profile F is an experimental-equilibrium if, for any agent A who deviates from its strategy in F by using another strategy from Σ, A does not increase its estimated expected utility, given that the other agents follow their strategies in F.*[8]

Definition 2 (Dominant Experimental Equilibrium) *A profile F is a dominant-experimental-equilibrium if, for any other profile F' that is an experimental-equilibrium, both the sellers and the buyers obtain the largest expected utility by following F.*

Our experiments are aimed at finding such equilibria.

6 Experiments

In order to reveal experimental equilibria, we have performed a series of experiments. Our main goal was to identify dominant experimental equilibria profiles, if these exist. Dominant equilibria profiles are desirable since, if they exist, all the agent that know them should prefer using the strategies of the dominant equilibrium. This will result in both utility maximization and stability. Hence, dominant equilibria profiles we find in our experiments should be recommended for the design of agents that buy and sell in fixed-price electronic markets similar to those we examine.

[8]Since our market is complicated, we assume that the set of strategies that is considered for deviation is determined in advance. Note that this definition differs from Nash equilibrium since we use an estimation of the expected utility, rather than the actual value.

Dimension / Setting	The utility function of the sellers			The utility function of the buyers		Stocks Size		The set of buyers		The size of the market	
	Risk Prone	Risk Neutral	Risk Averse	Conceding	Non - Conceding	Homo.	Hetero.	i.i.d	bi-polar	simple	larger
I		V		V		V		V		V	
II		V		V		V		V			V
III		V			V	V		V		V	
IV	V			V		V		V		V	
V			V	V		V		V		V	
VI		V		V			V	V		V	
VII		V		V		V			V	V	
VIII			V	V		V			V	V	

Figure 1: Experiments run with SEMI.

Our experiments were conducted using the SEMI model and simulator (described in Section 3). The settings of the experiments are summarized in Figure 1. Although we have experimented with several different market settings, some parameter values were not changed across settings, as follows. A seller's gain from selling one unit of the good is 1. A seller's cost for holding one unsold unit in stock for one iteration is 2 (i.e., $C_s = 2$). The underlying intuition for setting this cost is that, in similarity with real markets, it should be costly to hold items unsold in stock. Such costs result from the costs of both storage and financing of the unsold units. Determining the size of C_s was done experimentally. We have examined $C_s > 2$ as well, however $C_s = 2$ was found sufficiently large to express costliness of unsold stocks.

In cases I through VI (see Figure 1), in which the types of the buyers were stochastic independent and identically distributed (i.i.d.), the distribution of types was set to a normal distribution with mean $\mu = 50$ and standard deviation $\sigma = 40$. When smaller μ values (e.g., $\mu = 10$) were used in our experiments, some buyers ended up with negative types. One way to bypass this problem is by cutting the tail of the distribution. Yet, this will result in a non-normal distribution, which will, in turn, complicate the analysis of the results and impose difficulty in comparing them to other studies. We solved the problem by selecting a greater μ value, leaving virtually all buyer types positive.

In cases VII and VIII (Figure 1), where the types of the buyers were distinguishable, we devised a bi-polar distribution, which is in fact a superposition of two narrow normal distributions. The choice of a bi-polar distribution was made to model a common structure of B2B MRO markets, where buyers can usually be classified into larger-volume buyers and smaller-volume buyers, and the larger volumes are significantly larger than the smaller ones, with no overlap in size.

Across all simulations, each run consisted of 30 encounters. The number of encounters in our experiments should not be too small, since our strategies are applicable only for markets where buyers and sellers interact repeatedly. Yet, the number of encounter should not be too large either. This is because we consider short periods of time in which prices and stock sizes do not change, and long-term customer-vendor relationships do not form. Under such constraints, we opted for 30 encounters, gaining the additional advantage of statistically significant results.

In all simulation cases, in each encounter, for each buyer B^i and its type TY^i, the size of its purchase-order was chosen randomly from $\{TY^i - 1, TY^i, TY^i + 1\}$. Note that two agents of different types may have purchase-orders of equal sizes in a given encounter. Thus, the type of a buyer cannot be determined from its purchase-order (though it can be learned from multiple orders). The sizes of stocks held by buyers were computed by $(\mu - X_\sigma \cdot \sigma) \cdot |\mathcal{B}|/|\mathcal{S}|$, where X_σ is the factor by which we can constrain the stock to be smaller than the average expected cumulative demand μ, thus imposing supply shortages. $|\mathcal{B}|/|\mathcal{S}|$ is a normalization factor which is necessary to allow comparison between different market sizes.

Experiments were performed on two scales of markets: 1) a small market composed of 9 buyers and 3 sellers, and 2) large markets with 100 buyers trading with 14 sellers. A market of 9 buyers and 3 sellers is the smallest size in which meaningful many-to-many buyer-seller interactions occur. The number of sellers must be at least 3 so that a buyer will have the opportunity to select between at least two sellers, even in the case it decides not to return to one of the sellers. The number of buyers should be at least 3 times the number of sellers, to allow an average of at least 3 buyers per seller, so that a seller can choose from among them.

The exponential search space of possible deal combinations implies that even for the small market we cannot experiment with all possible combinations. Nevertheless, for the smaller case, running 100,000 experiments was sufficient to arrive at results which are statistically meaningful, and stable. In addition, for the small market it was possible to enumerate offline, prior to the experiments, the set of deal combinations to be experimented with. That is, for each simulation configuration we computed and stored set of $30 * |B| * 100,000$ deals (30 per buyer, for 100,000 runs), which were later used for the experiment. A stored set of deal combinations allows comparison between results of different experiments that use this stored set. For the larger market, however, the space complexity of the set of deal combinations is prohibitively large, thus pre-computation of the set is not an option. Therefore, we dynamically created the buyers' deals for each simulation run. The number of simulation runs for the large market was 400,000.

The basic setting that we considered is of risk-neutral homogeneous sellers and stochastic conceding buyers acting in a small market. In all of the other experiments we changed one parameter setting, keeping the other parameters as in the basic setting. First we changed the size of the market, i.e., we considered risk-neutral homogeneous sellers and stochastic conceding buyers acting in a *large market*. Second, we changed the utility function of the buyers and conducted experiments of risk-neutral homogeneous sellers and *non-conceding buyers* acting in a small market. Third we changed the risk attitude of the sellers and then we considered heterogeneous sellers. Finally, we considered non-stochastic buyers with both risk neutral and risk averse sellers. After presenting the results of these experiments, we will discuss our findings.

6.1 The basic setting: risk-neutral homogeneous sellers and stochastic buyers in small markets (case I)

In the basic setting the market consists of 9 conceding buyers and 3 risk-neutral sellers holding stocks of 100 units each. Experiments performed in this market show that the strategy profile (LoyalP RandS) is the dominant experimental equilibrium of these markets.[9] The results are shown in the central column of Figure 2. There, UB (US) is the average expected utility that the buyers (sellers) obtain by following the corresponding strategy.

In addition to the dominant equilibrium, we have identified 5 non-dominant experimental equilibria. Three of these include sellers that benefit from choosing the buyers according to the size of their orders (i.e., the sellers follow the OPO strategy). Sellers that implemented ORType, OType, DRType, DType with $C_t \in [0, 4]$ obtained poor expected utility. These strategies do not appear in the Figure as they are not part of any equilibrium. For these type revealing strategies, there was always an alternative greedy or uninformed strategy that yielded a higher remuneration.

6.2 Large markets (case II)

To find out whether the results obtained for the small market apply to larger markets, we have examined a market that consists of 100 buyers and 14 sellers. The size of each seller's stock was set to 233 units (in proportion to the 100 unit stock of the basic case with $X_\sigma = 0.412$). The results of these experiments show that the strategy profile (LoyalP, RandS) remains

[9]Significance of the results was tested with the t-tests with $\alpha = 0.05$.

	9 buyers 3 sellers		100 buyers 14 sellers
	Non−conceding buyers	Conceding buyers	Conceding buyers
B and S strategies	Stock=100	Stock=100	Stock=233
RandB,RandS	**V** UB: 0.5987 US: 0.8637	**V** UB: 0.6385 US: 0.8636	**XX**
RandB,OPO	**V** UB: 0.4867 US: 0.8637	**V** UB: 0.5334 US: 0.8637	**XX**
RandB,DPO	**V** UB: 0.4619 US: 0.8637	**V** UB: 0.6377 US: 0.8637	**X X**
LoyalP,RandS	**V** UB: 0.6602 US: 0.9758	**V** UB: 0.7087 US: 0.9854	**V** UB: 0.6846 US: 0.997
Loyal,RandS	**X**	**X**	**V** UB: 0.6847 US: 0.9969
Prob,RandS	**X**	**X**	**V** UB: 0.6845 US: 0.9914
Prob,OPO	**V** UB: 0.5153 US: 0.9653	**V** UB: 0.5673 US: 0.9653	**XX**
LoyalW,OPO	**V** UB: 0.5031 US: 0.952	**V** UB: 0.5545 US: 0.952	**XX**

The dominant experimental equilibrium

V An experimental equilibrium

X The corresponding strategies are not in experimental equilibrium

XX These strategies are undecidable. Nevertheless, (LoyalP RandS) dominates them.

Figure 2: Equilibria found for a market with risk-neutral sellers with homogeneous stocks and buyer types normally distributed with parameters $\mu = 50, \sigma = 40$.

the dominant experimental equilibrium. Numerical results are presented in the rightmost column of Figure 2. Note that, for some strategy profiles, no results regarding equilibria are presented. Our experiments have proven dominance of (LoyalP, RandS) over these profiles. Yet, a large market imposes a high time-complexity for checking whether a profile is in experimental equilibrium. Because of this complexity, we have refrained from checking experimental equilibria of all of the profiles which were dominated by (LoyalP, RandS). In particular, the profiles (LoyalW, RandS), (LoyalW, OPO), (Prob, OPO), (RandB, RandS), (RandB, OPO) and (RandB, DPO) were dominated by (LoyalP, RandS), hence we have not checked them for experimental equilibria.

As found in the small market, the strategy profiles (LoyalP, DPO), (LoyalP, OPO), (Loyal, DPO), (Loyal, OPO), (LoyalW, DPO) and (Prob,DPO) are not experimental equilibria thus are not presented in the Figure 2. Similarly, any strategy that required that the seller pay to acquire the type of the buyers was not part of any equilibrium. In difference from the results obtained for the small market, the profiles (Loyal, RandS) and (Prob, RandS) are experimental equilibria in the large market. Nevertheless, they are dominated by (LoyalP, RandS).

Altogether, the experimental results of this case strengthen our findings that (LoyalP, RandS) is the dominant experimental equilibrium in a large number of markets.

6.3 Non-conceding buyers (case III)

A buyer that follows the LoyalP strategy would return to a seller even if the seller only partially satisfied its order. In some cases, buyers are more sensitive to partial satisfaction of their orders. The displeasure of such non-conceding buyers with partial satisfaction is expressed by a significant reduction in their utility, as expressed in equation 4. We wanted to check whether in the case that buyers are non-conceding, the profile (LoyalP, RandS) will stop being the dominant experimental profile. It seems that the strategy LoyalW, according to which a buyer returns to a seller only if the seller completely satisfied its purchase-order, is more beneficial in this case. However, our experiments of risk-neutral homogeneous sellers and stochastic non-conceding buyers in small markets show that this is not the case. This is shown in the leftmost column in Figure 2, where (LoyalP, RandS) is the dominant experimental equilibrium.

6.4 Risk averse and risk prone sellers (cases IV and V)

Next we checked situations where the sellers are risk averse. It seems that it may be beneficial for a risk-averse agent to deviate from RandS where it makes decisions randomly. We considered situations of homogeneous sellers and stochastic conceding buyers in small markets. Here, we used the parameter Av to express the level of averseness of the seller. As Av decreases the averseness increases (see Section 3). We found that only when the sellers are highly averse to risk, (LoyalP, RandS) stops being a dominant strategy (see Figure 3). In particular, for $45 < Av < 100$ we found that the profile (LoyalP, RandS) is the dominant experimental equilibrium (as in markets with risk neutral sellers). For averseness values of 45 and 46, the expected utility of the sellers for the (LoyalP, RandS) profile was equal to their expected utility for the profile (Prob, OPO). Among these two profiles, the buyers' expected utility is greater for (LoyalP, RandS). Thus (LoyalP, RandS) is a weakly dominant profile. For $Av < 44$, i.e., sellers' averseness is intensified, their expected utility from the profile (Prob, OPO) is greater than their expected utility from (LoyalP, RandS). Nevertheless, a dominant experimental equilibrium was not found because the buyers benefit most when the sellers are indeed RandS.

We have further tested markets with risk-prone sellers. As expected, the dominant experimental equilibrium found for such markets is (LoyalP, RandS), as appears in Figure 3. Profiles that do not appear in the table are not in equilibrium.

		Buyers’ and sellers’ strategies, Stock=100, Ct=0					
Av	Sellers’ Risk	LoyalP, RandS	Prob, OPO	RandB,RandS	RandB,DPO	RandB,OPO	LoyalW,OPO
	Prone	US: 2.6489 UB:0.7088	US: 2.5214 UB:0.5674				
	Neutral	US: 0.9854 UB:0.7087	US: 0.9653 > UB:0.5673 <	US: 0.8636 UB:0.6385	US: 0.8637 UB:0.6377	US: 0.8637 UB:0.5334	US: 0.952 UB:0.5544
100	Averse	US: 6.9808 UB:0.7087	US: 6.9534 UB:0.5673				
46	Averse	US: 77.0135 = UB:0.7087 >	US: 77.0116 UB:0.5673				
44	Averse	US: 93.8886 < UB:0.7088 >	US: 93.8933 > UB:0.5674 <	US: 89.9745 UB:0.6386	US: 89.9786 UB:0.6378	US: 89.9753 UB:0.5334	US: 93.8583 UB:0.5544
35	Averse	US: 302.556 **X** UB:0.7087	US: 302.6356 > UB:0.5673 <	US:291.3932 UB:0.6387	US: 291.3886 UB0.6377	US: N/R UB:0.5334	US: N/R UB: 0.5544
20	Averse	US: 21998.7461 < UB:0.7086 >	US: 22002.6836 UB:0.5673				

(shaded box) The dominant experimental equilibrium

X This is not an experimental equilibrium

N/R Not Relevant, since U(B) of the buyers corresponding to these sellers is already smaller than the U(B) of the buyers who follow Prob when the sellers are OPO

Figure 3: Equilibria found for markets with three sellers behaving at different levels of averseness, nine conceding buyers with types normally distributed with parameters $\mu = 50, \sigma = 40$.

6.5 Heterogeneous stock sizes (case VI)

We hypothesize that in situations where sellers hold stocks of different sizes it will be beneficial for sellers holding a small stock to deviate from the RandS strategy; they may need to be more careful in satisfying their customers in order to attract them. However, experiments with heterogeneous sellers' stock sizes show that (LoyalP, RandS) is the dominant experimental equilibrium in this case too (see Figure 4). We considered a market of 3 sellers: one seller held a stock of 70 and two sellers held a stock of 100. Note that in this experiment we considered heterogeneous strategy profiles (i.e., sellers that hold stocks of different sizes implement different strategies). In addition, the utilities of the sellers were computed for each stock separately and not averaged over all sellers. That is, $U(S[70])$ in Figure 4 is the utility calculated for the seller with stock 70 and $U(S[100])$ is the average utility obtained by the sellers with stock 100.

We were able to find heterogeneous profiles that are experimental equilibria. For example, LoyalP for the buyers, OPO for the seller with stock 70 and RandS for the sellers with stock 100 is an equilibrium. Nevertheless, (LoyalP, RandS) is still the dominant experimental equilibrium (see Figure 4).

S0,S1, S2 and B strategies	Stock for S0, S1, S2 **70 100 100**
OPO,OPO,OPO LoyalW	**V** UB:0.4869 U(S[70]):0.9357 U(S[100]):0.9622
OPO,OPO,OPO Prob	**V** UB:0.4916 U(S[70]):0.9726 U(S[100])0.9723
OPO,RandS,RandS LoyalP	**V** UB:0.6187 U(S[70]):0.9939 U(S[100])0.9901
RandS,RandS,RandS LoyalP	**V** UB:0.6406 U(S[70]):0.9939 U(S[100])0.9906

Figure 4: Equilibria found for a market with $|\mathcal{B}|$=9,$|\mathcal{S}|$=3, sellers' stocks are $Stock(S^0) = 70$, $Stock(S^1) = Stock(S^2) = 100$ and buyer types are normally distributed with parameters $\mu = 50, \sigma = 40$.

6.6 Distinguishable buyers (cases VII and VIII)

One may hypothesize that a dominant equilibrium that includes RandS results from the buyer types being normally distributed. Thus, we decided to consider situations where there are only two possible types of buyers: low and high. In particular, we experimented with three sellers and a set of buyers composed of five buyers of type 20 and four buyers of type 70. The corresponding possible order sizes were 19,20,21 and 69,70,71 units.

The interesting result of this experiment is that (LoyalP, RandS) is no longer an experimental equilibrium. A single experimental equilibrium was found—(RandB, RandS)—which is, obviously, the dominant equilibrium. This result was obtained for both markets with risk-neutral and risk-averse sellers. It is somewhat surprising that, in the resulting equilibrium, RandB replaces LoyalP from preceding experiments. We would expect RandS to be replaced by a more deterministic strategy instead. Furthermore, this result implies that, in this market too, it is not beneficial for the sellers to learn the types of the buyers, even though these are clearly defined and distinguishable. We will discuss our intuition of this finding and the previous ones in the next section.

6.7 Discussion

The main result of our experiments is that the strategy profile (LoyalP, RandS) is the dominant experimental equilibrium in almost all of the situations considered. Our intuition is that this is the case because the (LoyalP, RandS) profile leads to the best stable split of the buyers among the sellers. To reach a good split, buyers completely unsatisfied by their current sellers should move to other sellers. When buyers who are completely unsatisfied implement a positive probability of staying with the non-satisfying sellers, as in the Loyal strategy, the resulting split of the market is non-beneficial. Note, however, that too frequent hopping to other sellers, either when buyers are partially satisfied (e.g., when they follow LoyalW) or when hopping is done regardless of the level of satisfaction (e.g., when RandB is followed), renders an unstable market. Instability of the market results in reduced utilities to both buyers and sellers. Thus, the profile (LoyalP, RandS) is the go between.

LoyalP behavior of buyers may be further understood as means for buyers to consider possible outcomes of future encounters in their current decision. For instance, a buyer which is partially satisfied by a seller at a given encounter may become fully satisfied in proceeding encounters, when agents that were not satisfied at all by that seller leave and the missing stocks for fully satisfying its order become available.

A seller's goal is to have a set of buyers that it will (almost) always be able to satisfy. Because of the dynamism implied by the buyers following LoyalP, a seller that holds a mixed set of customers, some of larger types and some smaller types, has a better ability to satisfy almost all. The best way for a seller to form such a mixed buyer set is to follow the RandS strategy, and this explains the sellers part in the (LoyalP, RandS) profile.

There were two cases where changes to the basic setting resulted in the profile (LoyalP, RandS) not being the dominant experimental equilibrium. The first was the case where risk attitudes were introduced. For sellers that were highly averse to risk, (LoyalP, RandS) was still an experimental equilibrium, however no longer the dominant one. Nevertheless, no other equilibrium became dominant. In this case, the sellers preferred (Prob, OPO) to (LoyalP, RandS); this is reasonable given their averseness to risk.

The second case was the case of distinguishable types. Under, the bi-polar distribution, where buyers were divided into two types, low and high, the profile (LoyalP, RandS) was no longer an equilibrium. Instead, the unique equilibrium was (RandB, RandS). Our intuition is that this stemmed from the specific choice of types in this experiment. There were 4 buyers of type 70 and 5 buyers of type 20. The average type (42.2) is lower than in the stochastic case (50) and the stock (100 per seller) was not changed. However, for the 4*70;5*20 type distribution, it is more difficult to split the market in a way that the buyers will be satisfied than it is in the case of normal distribution. Therefore, it is better for the buyers not to reach a stable split, but each time to try a different one by following RandB. When the buyers follow RandB, the sellers' strategy has no effect on the buyers' decision. Hence the best alternative for the sellers is to follow RandS.

When changes to the basic setting other than the bi-polar distribution were made, reaching a stable split was both possible and beneficial. For example, consider a non-conceding buyer. Although such a buyer loses more than a conceding buyer when partially satisfied, our intuition is that punishing the seller on partial satisfaction (e.g., using LoyalW) will not increase the expected utility of the buyer. Rather, as in the conceding case, when using LoyalW a stable split will not be reached. Thus, (LoyalP, RandS) is a dominant experimental equilibrium. Similarly, when there is one seller with stock 70 and two sellers of stock 100, the seller with the smaller stock can almost always sell all of its stock, regardless of the strategy that it uses. However, reaching a stable split is beneficial to the other large sellers. Therefore, (LoyalP, RandS) is the dominant equilibrium.

In addition to market settings discussed in this paper, we have experimented with, and analyzed, other settings. Results from experiments performed with markets where sellers held stocks of 70 and 150 units were analyzed and published in [5, 6]. There, it was found that for a stock size of 70, (LoyalP, RandS) was not a dominant equilibrium, although it was an equilibrium, and no other profile was a dominant equilibrium. We have no intuition that explains this behavior for this stock size. When the stock was 150, (Prob, RandS) was the dominant equilibrium, maintaining the RandS dominance, however replacing the buyers' strategy by a more punishing one, which is reasonable when the supply is increased.

In all of our experiments, for any setting considered, we observed that learning the type of the buyers is not beneficial to a seller. Our intuition is that this is because we considered situations where the demand was very close to supply. The sellers were quite often able to sell all their stock. In such cases it is better for sellers to form a set of mixed size customers, rather than spending money on identifying large buyers and trying to attract them. The best way to form such a set is using RandS.

It is important to note another property of RandS. When the order of arrival of buyers' requests at sellers' sites is random, and this arrival order does not depend on the quantity requested, using RandS is equivalent to using FIFS. This observation is particularly useful for sellers when the buyers can observe the order of their requests' arrival. The sellers can implement RandS since it is their dominant strategy, however the use of FIFS, when observable, exhibits a fair service. Since FIFS is equivalent to RandS, the sellers can implement FIFS with no harm to their gain.

7 Summary and Conclusions

The goal of this paper was to study strategic interactions among agents transacting in B2C and B2B MRO electronic markets, and to find stable and beneficial strategy profiles for such agents. In particular, strategies for buyers to select sellers and for sellers to decide which purchase-orders to satisfy were sought. Desirable criteria that recommended strategies should address are stability and maximization of gains. Since the profits of each agent participating in a trading activity are influenced by the other agents' activity, we sought strategy profiles that maximize each agent's utility *given* the actions taken by the other agents. The notion of equilibrium provides us with these desirable properties. A profile of strategies is in equilibrium if no agent can increase its utility by deviating to another strategy that is not part of the profile.

Traditionally, the study of strategic behavior of agents belongs to the field of game theory. Hence, our research is based on some game theoretic concepts. However, since practical electronic commerce scenarios are much more complex than classical games, our study includes several simplifications, as follows. We have limited the sets of strategies implemented by the agents to those relevant to the problem. In addition, since analytical computation of the expected utilities of the agents is very complex, we have opted for an approximation of these. Using a subset of the strategy space and approximated utilities, classical notions of equilibria do not hold. Hence, we defined the notion of experimental equilibrium. We empirically computed the expected utilities of the participating agents, and used these values to evaluate all of the combinations of strategies resulting from the sets we have assumed.

Based on our results, it is recommended that automated trading agents implement the strategies of the dominant experimental equilibria found. This will provide both buyers and sellers maximal gains. When the market includes stochastic, independent and identically distributed (i.i.d.) buyers—a typical buyer distribution in B2C electronic markets—it is recommended to implement the (LoyalP, RandS) profile. That is, sellers can be designed to select buyers at random, however buyers must be designed to be loyal to sellers as long as the latter satisfy at least part of their purchase orders. When the buyers' types are closer to the bi-polar distribution, the recommendation is to implement the profile (RandB, RandS). That is, both sellers and buyers should select one another in a random manner. It important to notice that these results are invariant to the sellers' attitude to risk[10], to the size of the stocks held by the sellers, to the buyers' attitude to partial satisfaction of their purchase orders, and to the size of the market. Since, in the dominant experimental equilibria found, sellers select buyers at random, learning the buyers' type cannot increase the sellers expected utility.

To summarize, we have developed a new means for studying strategic equilibria in complex, E-commerce environments where analytical computation of such equilibria is infeasible. We have specifically developed a model for B2C and B2B MRO electronic market, where prices are relatively uniform, buyers and sellers typically interact for short to medium periods of time, and the dynamism of the market results in occasional stock shortages. Subject to these market conditions and under a variety of settings, we have identified an array of strategic equilibria. In particular, we found dominant equilibria, and the strategies from which these equilibria are comprised are the ones recommended for use by trade agents, as they will provide them with both utility maximization and stability.

Electronic markets may include a huge number of buyers and sellers. The largest markets that we considered in this work consists of 100 buyers and 14 sellers. We believe that our ability to extend our results from small markets (3 sellers and 9 buyers) to the larger ones supports our hypothesis that these results will be also valid in larger markets as well. However, further research is needed to verify this hypothesis. Another issue that is problematic when considering the application of our results to real markets is our assumption that the set of strategies that is considered for deviation from an equilibrium is fixed. If additional strategies for the buyers and sellers would be identified, additional experiments should be carried to support the existence of the equilibria in the larger context.

In particular, in this work, we have studied agents' strategies that consider a history of only one time period. Future work should examine strategies that consider longer histories for decision making. Another aspect which deserves consideration in future research is dynamism. Agents may benefit by using adaptive strategies that enable them, in the course of interacting with other agents in the market, to alternate their behaviors. Whether longer

[10]This is true while the sellers are not too averse. Then, (LoyalP, RandS) is still in experimental equilibrium, but there is no dominant experimental equilibrium.

histories or adaptive behavior will provide more beneficial strategies is yet an open question to explore.

References

[1] L. F. Cranor and P. Resnick. Protocols for automated negotiations with buyer anonymity and seller reputations. *Netnomics*, 2(1), 2000.

[2] M. A. de Francesco. The emergence of customer markets in a dynamic buyer game. Quaderni del Dipartimento di Economia Politica 225, U. degli studi di Siena, 1998.

[3] E. Firedman and S. Shenker. Learning and implementation on the internet. Working Paper. http://fas-econ.rutgers.edu/home/friedman/research.htm, 1998.

[4] S. French. *Decision Theory: An Introduction to the Mathematics of Rationality*. Ellis Horwood Limited, 1986.

[5] C. V. Goldman, S. Kraus, and O. Shehory. Agent strategies: for sellers to satisfy purchase-orders, for buyers to select sellers. In Yves Demazeau and Francisco Garijo, editors, *Proceedings of the Tenth European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW01)*. 2001.

[6] C. V. Goldman, S. Kraus, and O. Shehory. Equilibria strategies for selecting sellers and satisfying buyers. *Cooperative Information Agents V*, LNAI 2182:166—177, September 2001.

[7] J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, May 2000.

[8] J. Sairamesh and J. O. Kephart. Price dynamics and quality in information markets. *Decision Support Systems*, 28(1-2):35–47, 2000.

[9] J. Sgall. On-line scheduling. *LNCS*, 1442:196–231, 1998.

[10] C. Shapiro. Consumer information, product quality, and seller reputation. *RAND Journal of Economics*, 13:20–35, 1982.

[11] D. R. Vincent. Modelling competitive behavior. *RAND Journal of Economics*, 23(4):590–599, 1992.

[12] B. Yu and M. P. Singh. Small-world reputation management in online communities. *CIA2000*, LNAI 1860:154–165, 2000.

[13] G. Zacharia. Collaborative reputation mechanisms for online communities. Master's thesis, MIT, 1999.

A Sellers' Strategies - An Example

Example 1 *Suppose a seller, S^j, has a stock of 14 units, i.e., $ST^j = 14$. In a given encounter k, four buyers, $B^1, ..., B^4$ approached S^j sequentially with the requests 4,5,8, and 4, respectively. These buyers' types are 3,6,9 and 5 respectively. The buyers will be supplied differently depending on the strategy used by the seller.*

RandS: *The distribution is not deterministic. For example, it is possible that B^1 and B^4 each, will obtain 4 and B^3 will obtain 6.*

OPO: *B^3 will get 8, B^2 will obtain 5 and either B^1 or B^4 will get 1.*

DPO: *B^1 will obtain 2 (i.e., $\lfloor \frac{4}{21} \cdot 14 \rfloor$), B^2 will obtain 3, B^3 will obtain 5 and B^4 will obtain 2. In addition, one of the agents that will be chosen randomly will obtain additional 2.*

ORType or OType *B^3 will obtain 8, B^2 will get 5, B^4 will obtain 1 and B^1 will not get anything. This division is similar to OPO, however, B^4 gets 1 in this case, while it may not get anything when OPO is used.*

DRType or DType: *B^1 will obtain 1 (i.e., $\lfloor \frac{3}{23} \cdot 14 \rfloor$), B^2 will obtain 3, B^3 will obtain 5, B^4 will obtain 3. In addition, one of the buyers will obtain an additional unit.*

As demonstrated in the example, a buyer may be supplied with different quantities depending on the strategy that is used by the seller. For example, B^1 may obtain 4 units when RandS is used, it obtains 2 when DPO is used, it obtains 1 when DType is used, it obtains nothing when OType is used and it may obtain 1 or nothing when OPO is used.

B Buyers' Strategies - An Example

Example 2 *Suppose there are three sellers S^1, S^2, and S^3, and in encounter k, B^i approached S^1 with a purchase-order $PO_1^{ik} = 5$. We consider three cases and indicate which seller B^i will approach at encounter $k + 1$.*

(1) S^1 supplied B^i 5 units (of 5) in encounter k:

If B^i uses the strategies Loyal, LoyalP, Prob or LoyalW it will approach S^1 in iteration $k+1$

as well. If B^1 uses the strategy RandB it will approach S^1 with a probability of $\frac{1}{3}$. With a probability of $\frac{1}{3}$ it will approach each of the other sellers.

(2) S^1 did not supply B^i anything in encounter k:
If B^i uses the strategies LoyalW, Prob and LoyalP it will not approach S^1 at encounter $k+1$. It will approach either S^2 or S^3, each with a probability of $\frac{1}{2}$. If B^i uses the strategies RandB and Loyal it approaches S^1 with a probability of $\frac{1}{3}$. With a probability of $\frac{1}{3}$ it approaches each of the other sellers.

(3) S^1 supplied B^i 3 units (of 5) in encounter k:
If B^i uses the strategies Loyal or LoyalP it approaches S^1 in encounter $k+1$. If B^i uses the strategy Prob it will approach S^1 with a probability of $\frac{3}{5}$. It will approach each of the other sellers with a probability of $\frac{1}{5}$. If B^1 uses the strategy RandB it will approach S^1 with a probability of $\frac{1}{3}$ and each of the other sellers with a probability of $\frac{1}{3}$.

If B^i uses the strategies LoyalW, it will not approach S^1 at encounter $k+1$. It will approach either S^2 or S^3, each with a probability of $\frac{1}{2}$.

As shown in the above example, a buyer B^i using the strategies Loyal, LoyalP, LoyalW and Prob demonstrates some degree of loyalty to the seller S^j, that it approached in the previous encounter. If S^j supplied all its request, B^i will return to it in the next encounter. These strategies differ in the cases of partial fulfillment and when S^j does not supply anything. Using the Loyal strategy the buyer presents the most loyal behavior. It returns to the seller if it was partially supplied, and chooses randomly between all the sellers (including S^j) when S^j supplied nothing. Using LoyalP, LoyalW and Prob, B^i does not return to S^j when it has not supplied it anything. These strategies differ in the case of partial supply. Using LoyalW, B^i does not approach S^j in such a case. Using LoyalP, B^i returns to S^j and using Prob, it approaches S^j with a probability that is proportional to its satisfaction level.