Equilibria Strategies for Selecting Sellers and Satisfying Buyers *

Claudia V. Goldman[1] **Sarit Kraus**[1,2] **Onn Shehory**[3]
clag@cs.biu.ac.il sarit@cs.biu.ac.il onn@il.ibm.com

[1] Department of Mathematics and Computer Science
Bar-Ilan University, Ramat-Gan 52900,Israel
[2] Institute for Advanced Computer Studies, University of Maryland
College Park,MD 20742
[3] IBM Research Lab in Haifa - the Tel-Aviv Site

Abstract. Dynamism of trade activity inevitably results in situations where sellers face local supply shortages. In such cases, sellers need to decide which buyer purchase requests to satisfy. Commonly, sellers satisfy purchase requests based on their arrival order, i.e., First In is First Served (FIFS). In electronic trade, sellers may follow strategies different from FIFS without the buyers being able to detect this difference. Buyers, in response to the sellers' strategic behavior, may themselves adopt strategies that will maximize their utility. Previous research has suggested strategies to be used by electronic seller-agents and buyer-agents. Yet, that research examined markets in which buyers are willing to accept partial satisfaction of their request and sellers' stocks are all the same. A simulation tool was developed under such conditions. This paper utilizes the simulation tool to explore equilibria in more realistic markets, where sellers' stocks are heterogeneous, and buyers suffer significant losses from partial satisfaction of their requests.

1 Introduction

Strategic behavior is inseparable from trade activity, and was implemented by traders for thousands of years. Electronic markets do not change this situation, as strategies can be devised and used to increase the gains of electronic trading parties too. However, in difference from traditional trade, electronic markets introduce a combination of conditions that does not exist in other markets. As a result, well-known strategies may become obsolete, and new, beneficial strategies will arise. This gives ground to research of the type presented in this paper.

Here, we examine scenarios where electronic sellers dynamically receive multiple purchase-orders. At times, the cumulative demand at a specific seller may surpass its stock level. Thus, the seller will not be able to address all orders, however, it would like to maximize gains nevertheless. We study strategies that allow this maximization. In business-to-consumer physical trade, sellers commonly satisfy purchase requests based on their arrival order, i.e., First In is First Served

* This research is supported in part by NSF under grant No. IIS9820657 and by IBM.

(FIFS). The FIFS strategy is considered fair, thus acceptable by buyers, and therefore sellers that follow it are not exposed to punishing buyer reactions. In electronic trade, sellers may follow strategies different from FIFS without the buyers being able to detect this difference. Such deviation from FIFS may increase the gains of the seller, due to preference of more profitable deals. Buyers, in response to the sellers' strategic behavior, may themselves adopt strategies that will maximize the chance of their orders being satisfied. Previous research [1] has suggested sets of strategies which are likely to be used by electronic seller-agents and buyer-agents in such trade conditions. Yet, that research examined markets in which buyers are willing to accept partial satisfaction of their request and sellers' stocks are all the same. While such conditions are excellent for developing a simulation tool (due to their simplicity), they do not hold in real markets. This paper utilizes the simulation tool of [1] to explore equilibria in more realistic markets, where sellers' stocks are heterogeneous, and buyers suffer significant losses from partial satisfaction of their requests. In particular, given these conditions, we study equilibria strategies that allow sellers to maximize gains and buyers to select sellers given the sellers' purchase-order satisfaction.

Our problem domain refers to multiple, possibly anonymous, buyers that electronically approach multiple automated sellers' sites. Buyers have no information regarding other buyers who visit the same sites, nor do they know in advance what is the stock level held by the sellers they approach. Sellers in our domain are uncertain regarding the number of buyers that will approach their site and the volume of their purchase-orders. The dynamism of electronic trade intensifies this problem. Given these settings, we devise strategies for sellers and counter-strategies for buyers, such that sellers maximize gains, buyers maximize satisfaction, and the market is in equilibrium. Due to the complexity of the problem, solving it analytically is practically infeasible, hence we opted for a simulation-based solution. The details are presented in the subsequent sections.

2 The SEMI system

The Simulation of Electronic Market Interactions (SEMI) system supports simulation of repeated encounters between sets of automated buyers and sellers. Although SEMI can be used for a broad range of settings, we use it to simulate sellers that hold limited stocks and buyers that place a purchase order with one seller at each encounter. In this study we are mainly concerned with markets in which, at each encounter, a seller may receive requests for goods that cumulatively exceed the amount available in its stock. Therefore, it needs to decide which of the requests to satisfy and which buyers to leave unsatisfied. Electronic sellers that face such decision problems should be equipped with strategies for making such decisions, or means for locating such strategies. Preferably, strategies followed by market players (sellers and buyers) should maximize gains and be in equilibrium. Unfortunately, analytical computation of such strategies is too complex for the market settings we study. Running simulations is important when analytical solutions are difficult to obtain. Experiments with SEMI can (and do) unravel equilibria between the buyers and sellers' strategies. These equilibria are the basis for recommendations for strategy design principles for automated buyers and sellers. SEMI was developed as part of an earlier study [1].

A SEMI simulation consists of a sequence of K encounters between sellers and buyers from the finite sets $\mathcal{S}$ and $\mathcal{B}$ respectively. In each encounter, k, a seller $S^j \in \mathcal{S}$ holds a stock of goods ST^{jk} for sale. Each buyer $B^i \in \mathcal{B}$ is associated with a type TY^i and a purchase-order PO^{ik}. TY^i indicates the average purchase-order size of B^i. B^i submits its purchase-order PO^{ik} to one of the sellers, S^j. S^j may sell B^i at most the amount specified in PO^{ik}. The utility of B^i increases in proportion to the portion of PO^{ik} supplied by S^j. A seller's utility increases with the cumulative amount that it sells to the buyers, and decreases with the amount of goods that is left in its stock at the end of the encounter.

We distinguish between two kinds of buyers: a *recognizable* and an *unrecognizable buyer*. A seller can use the information it obtained about the type of a recognizable buyer in future encounters. Information about the type of an unrecognizable buyer can only be used in the encounter in which it was obtained.

Here, we analyze the sellers and buyers agents' behaviors in three scenarios implemented with SEMI. All the sellers sell the same type of a good, where quality, price and units are the same as well. Thus, PO^{ik} specifies the number of whole units of the good that B^i would like to buy in encounter k. Even though sellers and buyers transact repeatedly, they do so for a rather short period, during which it is reasonable to assume that the price of a sold unit remains static.[1] We also assume that the prices are uniform; we do not handle auctions but transactions performed in fixed-price, catalogue-based, electronic stores. We further assume that each buyer is associated with only one purchase-order and each purchase-order is valid for only one encounter. That is, when a buyer B^i approaches a seller S^j with PO^{ik}, S^j can at most supply PO^{ik} in encounter k. It cannot delay the supply, and PO^{ik} cannot be supplied by another seller at any encounter. After B^i approached S^j at encounter k, we denote its order PO^{ik}_j.

In all of the experiments performed, a seller gains from the sale of each unit and it incurs a cost for each unit it holds unsold in its stock at the end of each encounter. In cases where a seller reveals buyers' types, it has an additional cost for doing this. The overall utility of a seller is the sum of its utility from each encounter. Formally, the utility of S^j in a given encounter k, when it receives r purchase-orders, is $U^j(PO^{1k}_j, \ldots, PO^{rk}_j)$ as described in table 1.

In the basic scenario, we assume that the size of the stock of all sellers at the beginning of each encounter is the same. A buyer B^i's utility is $U^i(PO^{ik})$ as presented in table 1. In the second scenario, we study non-conceding buyers. The utility from partial satisfaction of purchase-orders has a greater impact on non-conceding buyers. A non-conceding buyer B^i's utility is given by the function U^i_{nc} as described in table 1, where LS is the factor that reflects the level of buyer dissatisfaction with partial purchase-order fulfillment. We say that buyers who implement such utility function are non-conceding, since their willingness to accept partial purchase-order fulfillment is lesser than it was in the basic scenario.

[1] See reference to work by Kephart as explained in Section 6.

Note that here, the utility of a partially satisfied buyer may be nullified.[2] As the value of LS increases, the buyer's level of non-concession increases too. In the

Notation	Description
k	The k^{th} encounter.
K	The number of encounters.
$\mathcal{B}=\{B^i\}, 1 \le i < \infty$	The set of buyers.
$\mathcal{S}=\{S^j\}, 1 \le j < \infty$	The set of sellers.
ST^{jk}	The stock that seller S^j holds at the beginning of encounter k.
PO_j^{ik}	The purchase-order placed by B^i to S^j at encounter k.
$sat(PO_j^{ik})$	The actual deal satisfied by S^j after 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 good.
r^*	The number of times the seller bought information about the buyers' types in the given encounter.
$U^j(PO_j^{1k},\ldots,PO_j^{rk})$	$U^j(PO_j^{1k},\ldots,PO_j^{rk}) = G \cdot \sum_{i=1}^{r} sat(PO_j^{ik}) - C_s \cdot (ST^{jk} - \sum_{i=1}^{r} sat(PO_j^{ik})) - C_t \cdot r^*$
$U^i(PO^{ik})$	$U^i(PO^{ik}) = sat(PO^{ik})/PO^{ik}$
$U^i_{nc}(PO^{ik})$	$U^i_{nc}(PO^{ik}) = (sat(PO_j^{ik}) - (PO_j^{ik} - sat(PO_j^{ik})) * LS)/PO_j^{ik}$

Table 1. Summary of the SEMI's Notations

third scenario, we consider heterogeneous markets (i.e., markets with different stocks' sizes). The buyers' utility is given by $U^i(PO^{ik})$ as in the basic setting. The sellers may hold stocks with different sizes.

Since we cannot provide a full set of strategies because many variations of strategies exist, we consider representative sets of strategies for the agents.

2.1 The Sellers Agents' Strategies (Σ^s)

A seller's strategy specifies which portion of each purchase-order received to supply. The decision that stems from such a strategy could be based on: (i) the arrival time of the purchase-order (within the encounter); (ii) the size of the purchase-order; and (iii) the type of the buyer that submitted the purchase-order, when these buyers are recognizable.

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 are not cumulatively larger than its stock, the differences between the strategies become unimportant, as the seller agent will supply all the requests, regardless of the specific strategy.

[2] 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 equal to not being satisfied at all.

Uninformed seller: The seller does not consider the size of the purchase-orders nor the buyers' types (issues (ii) and (iii)).

1. RandS (Random Seller) — A seller S^j chooses randomly which purchase-order PO_j^{ik} to fulfill constrained by its stock. Assuming that the order of arrival of the buyers to the sellers does not depend on any characteristic of the buyers, this behavior is equivalent to a FIFS behavior.

Greedy Seller: This seller uses the size of the purchase-orders (ii).

1. OPO (Ordered Purchase-Orders) — The order of fulfillment of the buyers' purchase-orders is in decreasing order of their size.
2. DPO (Distributed Purchase-Orders) — A seller satisfies the purchase-orders proportionally to their size. If S^j receives r purchase-orders in encounter k, then each buyer B^i will be supplied with $\lfloor (PO_j^{ik} / \sum_{l=1}^{r} PO^{lk}) \cdot ST^j \rfloor$. The remainder of this distribution is allocated to one buyer, selected randomly.

Intelligent seller: This seller uses the buyers' types in its decision (iii).

1. ORType or OType (Ordered [Recognized] Types) — The order of fulfillment of the buyers' purchase-orders is in decreasing order of the type of the buyers. R in the prefix denotes recognizable buyers. When the buyers are unrecognizable (no R), the seller needs to pay for the buyer's type each time it would like to use it.
2. DRType or DType (Distributed [Recognized] Types) — A seller satisfies the purchase-orders proportionally to the buyers' types. That is, if S^j receives r purchase-orders in encounter k, then for each buyer B^i, the seller computes $po_i = \lfloor (TY^i / \sum_{l=1}^{r} TY^l) \cdot ST^j \rfloor$. If $po_i \geq PO_j^{ik}$ then po_i is set to PO_j^{ik}. The remainder is allocated as in DPO. The prefix R settings of ORType and OType hold here as well.

We hypothesize that a seller will benefit from satisfying highly typed buyers so that the latter return to them, and increase their gain. Moreover, we hypothesize that if a seller assumes that an unsatisfied buyer will limit the purchase-orders that it sends to a seller which have disappointed it, the seller should try to satisfy these buyers as much as it can. Even when the type is unknown, sellers can use the size of the buyer's current purchase-order to estimate its type.

2.2 The Buyers Agents' Strategies (Σ^b)

In the market set in our paper, at each encounter, a buyer only knows what portion of its 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 last encounter.[3]

[3] 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. Here, we focus on short-term interactions.

1. RandB (Random Buyer) — The buyer chooses a seller randomly.
2. Loyal — If seller S^j has satisfied buyer B^i *completely* or *partially* at encounter k, then B^i returns to seller S^j at encounter $k+1$. If S^j has supplied B^i nothing at encounter k, then at encounter $k+1$ B^i approaches a seller from $\mathcal{S}$ chosen randomly (it may approach S^j again).
3. LoyalP (Loyal and Punish) — If seller S^j has satisfied buyer B^i *completely* or *partially* at encounter k, B^i returns to seller S^j at encounter $k+1$. If S^j has not supplied B^i anything at encounter k, then at encounter $k+1$ B^i approaches a seller from $\mathcal{S} \setminus \{S^j\}$ chosen randomly.
4. LoyalW (Loyal Weak) — If seller S^j has satisfied buyer B^i *completely* at encounter k, then B^i returns to seller S^j at encounter $k+1$. Otherwise, B^i randomly chooses from among the other sellers ($\mathcal{S} \setminus \{S^j\}$).
5. Prob (Probabilistic Buyer) — If buyer B^i has approached seller S^j at encounter k with a purchase order PO_j^{ik}, then B^i will approach S^j at encounter $k+1$, with a probability of $sat(PO_j^{ik})/PO_j^{ik}$. This probability expresses actual satisfaction of a buyer with respect to its purchase order.

The buyers will always return to a seller that has completely satisfied them. These strategies differ when the buyers are partially satisfied or not satisfied at all. A buyer *punishes* a seller it has approached at encounter k by not returning to this seller at encounter $k+1$. LoyalW is the most punishing strategy. The probabilistic strategy is less punishing than LoyalW since a buyer may return to a partial satisfying seller with a positive probability.

Our hypothesis was that not returning to a partially satisfying seller will be the best strategy for the buyer. However, our results do not always support this.

3 Evaluation Criteria

The utility of each agent participating in an E-market is strongly influenced by the decisions taken by each one of the agents. Since in E-markets as those that are studied here, agents are self-interested and competitive, there is no single evaluation criterion that fully captures the preferences of all the market participants. Some solutions might be optimal for part of the agents, but disadvantageous for others. Our study seeks solutions, in which all of the agents get the maximal benefit given the other agents' behavior, and therefore they will have no incentive to deviate from these solutions.

We assume 2 sets of strategies, Σ^s and Σ^b, are given. A strategy profile F is a sequence of strategies, one for each buyer and one for each seller from the relevant sets of strategies. Given a profile F we will use SEMI to compute the average utility of each agent and use it as an estimation of its expected utility. Then, we look for combinations of strategy profiles for the sellers and buyers agents that are in *experimental equilibrium* [1]:

Definition 1 (Experimental Equilibrium). *We say that 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.[4]

Definition 2 (Dominant Experimental Equilibrium). *We say that 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.*

We define the concept of level of liquidity of the buyers to analyze the results.

Definition 3 (Level of Liquidity of the Buyers). *The level of liquidity of the buyers is defined by the average number of times that a buyer changes a seller during one SEMI encounter. Formally, it is given by $\frac{\Sigma_{|Sim|,K,|\mathcal{B}|} SellerChanges}{|Sim| \cdot K \cdot |\mathcal{B}|}$, where $|Sim|$ denotes the number of simulations run, and SellerChanges are the number of times that one buyer has changed a seller during one encounter.*

We hypothesize that a lower level of liquidity of the buyers results in a higher expected utility of both buyers and sellers, as long as the buyers change sellers at least once. Our results support this hypothesis.

4 Experiments with the Basic SEMI System

The experiments performed in the basic setting were extensively reported in our previous work [1]. Here, we summarize these findings to compare them with the newer results found in Section 5. We have chosen to report on experiments run on a market that consists of 9 buyers and 3 sellers. Although we have examined other market sizes, this size appears to be the smallest 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 decided 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 settings of the parameters used for the basic experiments are presented in Table 2. Note that two agents of different types may have the same purchase-order 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). We tested the behaviors of the agents for three values of stocks. The basic case was set so that some purchase-orders are always satisfied and that sellers are not always able to sell their entire stock. The results presented below are averages over 100,000 runs. In each run we used the same parameter settings but randomly chose purchase-orders for each buyer in each of the 30 encounters.

[4] 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.

$\lvert\mathcal{B}\rvert = 9$, $\lvert\mathcal{S}\rvert = 3$
$G = 1$, a seller's gain from selling one unit.
Types were normally distributed with $\mu = 50$ and $\sigma = 40$.
$PO_j^{ik} = \{TY^i - 1, TY^i, TY^i + 1\}$
for any k *and* $S^j, S^{j'} \in \mathcal{S}$, $ST^{jk} = ST^{j'k}$
$C_s = 2$, a seller's cost for holding an unsold item.
$ST^{jk} = (\mu - X_\sigma \cdot \sigma) \cdot \lvert\mathcal{B}\rvert / \lvert\mathcal{S}\rvert$, X_σ constrains the stock to be smaller than the average, $\lvert\mathcal{B}\rvert/\lvert\mathcal{S}\rvert$ is the normalization factor.

Table 2. Basic Settings

We report on all the experimental equilibria profiles found for stocks of 70, 100 and 150 units.[5] The results appear in Figure 1, and their interpretation follows. In the Figure, the numbers stand for the average expected utility the buyers (UB) and the sellers (US) obtain when following the corresponding strategy.

B and S strategies	Stock=70	Stock=100	Stock=150
RandB,RandS	V UB: 0.4823 US: 0.915	V UB: 0.6385 US: 0.8636	V UB: 0.8214 US: 0.7604
RandB,OPO	V UB: 0.3936 US: 0.9149	V UB: 0.5334 US: 0.8637	V UB: 0.7841 US: 0.7602
RandB,DPO	V UB: 0.48 US: 0.9151	V UB: 0.6377 US: 0.8637	V UB: 0.8244 US: 0.7604
LoyalP,RandS	V UB: 0.5029 US: 0.9963	V UB: 0.7087 US: 0.9854	V UB: 0.9217 US: 0.8669
LoyalP,OPO	V UB: 0.3556 US: 0.9965	X	X
Prob,RandS	X	X	V UB: 0.9285 US: 0.8703
Prob,OPO	X	V UB: 0.5673 US: 0.9653	X
LoyalW,OPO	X	V UB: 0.5545 US: 0.952	X

The dominant experimental equilibrium
V An experimental equilibrium
X The corresponding strategies are not in experimental equilibrium

Fig. 1. Equilibria found for a market with $\lvert\mathcal{B}\rvert$=9,$\lvert\mathcal{S}\rvert$=3, and types normally distributed with parameters $\mu = 50, \sigma = 40$.

Notice that, for larger stocks, the sellers' expected utility decreases, since more units remain unsold, and the sellers pay for this. In contrast, the buyers' expected utility increases since the chances to be fully satisfied are larger.

Dominant experimental equilibria exist for [100] and [150]. In both cases, the sellers benefit most by implementing the RandS strategy. For [100], there are five additional experimental equilibria, three of which include sellers that benefit by choosing the buyers based on their deals (i.e., the sellers follow the OPO strategy). Sellers that implemented ORType, OType, DRType, DType

[5] For brevity, we annotate these settings by [70], [100] and [150], respectively.

with $C_t \in [0,4]$ obtained poor expected utility; in all the cases there was a greedy or uninformed behavior that yielded a higher remuneration. For [150], there are four additional experimental equilibria.

For [70], five experimental equilibria were found, but neither dominates the others. Sellers that follow the OPO strategy gain a slightly (though significant) higher expected utility than sellers that implement RandS. Sellers that hold small stocks benefit from making buyers that behave LoyalP stay with them

When the buyers play RandB, the average utility of the sellers does not depend on their strategy. This can be explained by the fact that the sellers try to influence the choice of the buyers by their behavior. However, if a buyer chooses a seller randomly, regardless of the sellers behavior, the exact strategy that a seller uses is irrelevant. The three combinations of strategies comprising RandB buyers, and sellers who can act RandS, OPO, and DPO are experimental equilibria in the three settings of our simulations. Given that all of the buyers behave RandB, no single buyer will be motivated to deviate. Still, this equilibria is not dominant because, for the different stock sizes tested, there is always a punishing strategy for the buyers that yields higher expected utility for the buyers than when the buyers play RandB.

Interpretation of the Results Since the dominant experimental equilibria include random sellers, learning the buyers' types cannot increase the sellers expected utility. We discuss this issue further in Section 5.

When a dominant experimental equilibrium is found, sellers that implement RandS conduce the market to a low and adequate level of liquidity. Sellers that act RandS also avoid arriving at a static local minima as it is achieved by the sellers that follow DPO.[6] As opposed to DPO, OPO sellers may cause the buyers to change sellers even when their distribution was good.

When the stock size is rather small, it seems that the appropriate recommendation for the sellers is to implement the OPO strategy. Since the stock is small, the seller will benefit most by not loosing the buyers he can satisfy, and should not take risks by following the RandS strategy.

As explained in Section 2.1, our implementation of RandS resembles the FIFS strategy, when the buyers approach a certain seller in a random order. FIFS is commonly used in physical stores. There, since buyers can see the arrival order, FIFS provides a sense of fairness. In virtual stores, though, a buyer is usually not aware of the other buyers. As resulted in our simulations, FIFS is a beneficial strategy for automated sellers as well when they hold equally-sized stocks, and the order of the arrival of the buyers to a given seller is random.

Our results show that the buyers should follow the strategy that punishes most the sellers, given their stock size. As the stock size becomes larger, the competition among the buyers decrease, and therefore the buyers can more severely

[6] DPO sellers lead to a market with the lowest possible level of liquidity because DPO implies that the majority of the buyers are at least partly satisfied, so they do not switch to other sellers. This is also true for "bad" situations, those in which buyers and sellers could benefit from the change.

punish the sellers who do not satisfy them. All of the equilibria found included a punishing strategy for the buyers that yielded a higher expected utility than a less-punishing strategy.

5 Experiments with non-Conceding Buyers and Heterogeneous Markets

Non-conceding Buyers In real markets, sellers seem to consider the type of the buyers for decisions regarding serving them. The results presented in Section 4 contrast this expectation, as sellers maximize benefits by random, type independent, choice of the buyers to be served. One may hypothesize that this resulted from utility functions implemented in the basic SEMI setting. These functions imply that a buyer's utility is proportional to the level to which his/her purchase-order was satisfied.

To examine this claim, we ran another set of simulations, where non-conceding buyers react to partial satisfaction of their purchase-order in an amplified dissatisfaction, expressed via a reduced utility. The utility of a buyer B^i was changed to be $U^i_{nc}(PO^{ik})$ as explained in Section 2.

The equilibria found for a market with 9 buyers, 3 sellers, a stock of 100 units and LS=0.5 are exactly the same equilibria found for the basic setting (see Figure 2). Notice that in all of the cases, the expected utility of the buyers is smaller than the one obtained in the basic setting due to the LS factor influence. Moreover, there exists a dominant experimental equilibria that is achieved when the buyers implement the LoyalP strategy and the sellers, the RandS strategy.

B and S strategies	Stock=100
RandB,RandS	V UB: 0.5987 US: 0.8637
RandB,OPO	V UB: 0.4867 US: 0.8637
RandB,DPO	V UB: 0.4619 US: 0.8637
LoyalP,RandS	V UB: 0.6602 US: 0.9758
Prob,OPO	V UB: 0.5153 US: 0.9653
LoyalW,OPO	V UB: 0.5031 US: 0.952

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

Fig. 2. Equilibria found for a market with $|\mathcal{B}|$=9,$|\mathcal{S}|$=3, and types normally distributed with parameters $\mu = 50, \sigma = 40$. The leftmost figure considers non-conceding buyers. The rightmost figure considers sellers whose stocks are $Stock(S^0) = 70$, $Stock(S^1) = Stock(S^2) = 100$

The level of liquidity of the non-conceding buyers is larger than the level found in the basic setting. Even when the buyers' strategy induces them to return to a seller that has partially satisfied them, the buyers may behave as if the sellers have not satisfied them at all. This change of behavior applies to strategies that allow buyers to return to a partially satisfying seller.

Heterogeneous Markets As the results of the previous experiments suggest, when sellers are equal, sellers will maximize benefits by following RandS. Yet, when the sellers hold different stocks, competition among the sellers can arise such that acquiring the buyers' types may help the sellers. We have run another set of simulations for a market with 9 buyers whose expected utility is $U^i(PO^{ik})$ as in the basic setting, and 3 sellers that hold different stocks to examine this hypothesis. Without loss of generality, seller S^0 holds a 70 unit stock, and sellers S^1 and S^2 hold each a stock of 100 units. The results are shown in Figure 2. Note that the utilities of the sellers were computed for each stock, i.e., $U(S[70])$ is the utility calculated for S^0 and $U(S[100])$, the average utility obtained by S^1 and S^2).

In the basic setting, RandB buyers were always part of an experimental equilibria. The results of the the heterogeneous experiments are different. Buyers are better-off deviating from RandB to LoyalP regardless of the sellers' strategies. The increased competition among the sellers enables the buyers to punish the sellers more severely. As can be seen from Figure 2, the combination of LoyalP buyers and RandS sellers is still the dominant experimental equilibria as it was in the basic setting for an homogeneous market with [100]. Note that, except for equilibria that include RandB buyers in the basic case (for [100] and [70]), the experimental equilibria arrived at in the heterogeneous case are the same as those arrived at in the basic case.

S^0 benefits less in the heterogeneous market than it does when all sellers hold stock 70 ($U(S[70])$=0.9939 in the heterogeneous market, and $U(S[70])$=0.9963 in the homogeneous market). This is due to the fact that larger sellers can satisfy more buyers' deals than S^0 who holds a smaller stock. These LoyalP buyers will keep returning to S^1 or S^2 and will not approach S^0 as long as they are at least partially satisfied. In contrast, S^1 and S^2 benefit from the competition ($U(S[100])$=0.9906 in the heterogeneous market, and $U(S[100])$=0.9854 in the homogeneous market). The buyers' benefit from the heterogeneous market is greater than their benefit in a homogeneous market where all the sellers hold 70 units though smaller than their benefit in a homogeneous market with sellers holding each 100 units.

Notice that the FIFS strategy, which was the dominant experimental equilibrium in the basic case, is still in equilibrium. And, the experiments suggest no evidence that seller diversity effects the need for type information.

6 Related work

Scheduling algorithms (e.g., [3]) cannot be applied to E-markets because there is no global function which all the agents are trying to minimize. Agents react to each other's behaviors. Game theory and economics research have addressed issues of competition among buyers and of sellers that need to choose which buyer to supply (e.g., [4]). Nevertheless our assumptions differ from those assumed in these works. In our framework, buyers may consider the sellers' reputation when they choose which seller to approach. Reputation was studied as a function (e.g., [6]), as a social mechanism for avoiding untrustworthy parties [5], and as an expectation for quality. In our case, the buyers' strategy may change according to the service received from the sellers. The behavior of automated sellers situated in E-markets was investigated by Kephart et al. [2]. In particular, their focus was on dynamic changes in prices. There, time range was longer, and variations in stock were not handled as we do in this paper.

7 Conclusion

We examined strategies that seller agents can use to select buyers' purchase-orders in the face of limited stocks. We studied strategies that buyer agents can use to select sellers, based on their level of satisfaction.

It is unlikely to assume that both buyers and sellers will act as symmetric groups in E-markets. But, surprisingly, dominant experimental equilibria were found in our experiments, that show such symmetry. This suggests that the strategies' profiles that are in such equilibria are advantageous, and designers of automated electronic sellers and buyers should implement them. Agents will only loose utility by deviating from the group behavior.

We have found all of the experimental equilibria for the three scenarios of markets tested when all the sellers hold equally-sized stocks, variable-sized stocks and when the buyers are non-conceding. The main conclusion of this work is that sellers should behave randomly. RandS leads to a lower level of liquidity of the buyers in the market. As we hypothesize and show, lower levels of liquidity of buyers usually entail increased utility. Nevertheless, designers of agents for E-markets should be aware that local minima can be arrived at. The recommendation for the buyers is to punish the non-satisfying sellers as much as the stock size enables them, i.e., the severity of the punishment is inversely proportional to the amount of competition among the buyers.

References

1. C. V. Goldman, S. Kraus, and O. Shehory. Agent strategies: for sellers to satisfy purchase-orders, for buyers to select sellers. In *Proceedings of MAAMAW01*. 2001.
2. J. O. Kephart, J. E. Hanson, and A. R. Greenwald. Dynamic pricing by software agents. *Computer Networks*, 32(6):731–752, May 2000.
3. J. Sgall. On-line scheduling. *LNCS*, 1442:196–231, 1998.
4. D. R. Vincent. Modelling competitive behavior. *RAND Journal of Economics*, 23(4):590–599, 1992.
5. B. Yu and M. P. Singh. Small-world reputation management in online communities. *CIA2000*, LNAI 1860:154–165, 2000.
6. G. Zacharia. Collaborative reputation mechanisms for online communities. Master's thesis, MIT, 1999.