ECAI 2010
H. Coelho et al. (Eds.)
IOS Press, 2010

doi:10.3233/978-1-60750-606-5-575

An Equilibrium Analysis of Competing Double Auction Marketplaces Using Fictitious Play

Bing Shi and **Enrico H. Gerding** and **Perukrishnen Vytelingum** and **Nicholas R. Jennings** [1]

Abstract. In this paper, we analyse how traders select marketplaces and bid in a setting with multiple competing marketplaces. Specifically, we use a fictitious play algorithm to analyse the traders' equilibrium strategies for market selection and bidding when their types are continuous. To achieve this, we first analyse traders' equilibrium bidding strategies in a single marketplace and find that they shade their offers in equilibrium and the degree to which they do this depends on the amount and types of fees that are charged by the marketplace. Building on this, we then analyse equilibrium strategies for traders in competing marketplaces in two particular cases. In the first, we assume that traders can only select one marketplace at a time. For this, we show that, in equilibrium, all traders who choose one of the marketplaces eventually converge to the same one. In the second case, we allow buyers to participate in multiple marketplaces at a time, while sellers can only select one marketplace. For this, we show that sellers eventually distribute in different marketplaces in equilibrium and that buyers shade less and sellers shade more in the equilibrium bidding strategy (since sellers have more market power than buyers).

1 Introduction

Exchanges, in which securities, futures, stocks and commodities can be traded, are becoming ever more prevalent. Now, many of these adopt the double auction market mechanism which is a particular type of two-sided market with multiple buyers (one side) and multiple sellers (the other side). Specifically, in such a mechanism, traders can submit offers at any time in a specified trading round and can be matched by the marketplace at a specified time. The advantages of this mechanism are that traders can enter the marketplace at any time and they can trade multiple homogeneous or heterogeneous items in one place without travelling around several marketplaces. In addition, this mechanism provides high allocative efficiency [4]. These benefits have led many electronic marketplaces to also use this format. For example, Google offers DoubleClick Ad Exchange, which is a real-time double auction marketplace enabling large online ad publishers, on one side, and ad networks and agencies, on the other side, to buy and sell advertising space. However, because of the globalised economy, these marketplaces do not exist in isolation. Thus they compete against each others to attract traders and make profits. For example, stock exchanges compete to attract companies to list their stocks in their marketplaces and Google's DoubleClick Ad Exchange competes against other ad exchanges, such as Microsoft's AdECN and Yahoo!'s Right Media. Thus such competition is becoming an increasingly important area of research. Specifically, for such contexts, there exist two key research challenges. The first is how traders behave in the competing marketplaces environment, which includes how traders select marketplaces and how they bid in them. Then given the traders' behaviour for selecting marketplaces and bidding, the second issue is how competing marketplaces should set their fees to maximise their profits, while at the same time maintaining market share at a good level. In this paper, we address the first problem and use fictitious play to analyse the trader's equilibrium behaviour when a trader can choose to trade in multiple double auction marketplaces. We also analyse how market fees can affect this equilibrium. The insights from this analysis will be useful to guide a marketplace to set its fees.

A number of existing empirical and theoretical works analyse competing marketplaces. In this context, an annual Market Design Competition (CAT) was introduced as part of the Trading Agent Competition (TAC) [5]. Here, entrants need to design effective market policies and set appropriate fees to attract traders and so make profits. However, work in this area is still largely empirical in nature. On the other hand, a number of theoretical models have been proposed to analyse two-sided competing marketplaces (e.g. [1, 6, 7]). However, these works do not consider auction mechanisms to match traders and set transaction prices. Instead, they assume that traders only select marketplaces based on the *number* of other traders in the marketplace. In doing so, they assume that all traders are homogeneous (i.e. have the same preferences), and the marketplace has complete information about the preferences (also called the types) of traders. In real-world auction marketplaces, however, traders are usually heterogeneous and they are likely to have privately known preferences. Moreover, transaction prices are usually set according to the marketplace's pricing policy, which is affected by current demand and supply. To tackle these limitations, in our previous work [2], we proposed the first game-theoretic framework to analyse the competing double auction marketplaces, in which traders have different types (i.e. preferences). However, in this work we introduced a number of simplifying assumptions. First, we only considered a restricted number of discrete types. Second, we assumed that traders bid truthfully. However, in real-world marketplaces, traders often shade their offers in order to extract more profits from transactions. Furthermore, we only considered traders who can participate in a single marketplace at a time (so-called *single-home trading*), in contrast to the more generic *multi-home trading* where traders can participate in multiple marketplaces at a time.

In this paper, we address all of the above-mentioned shortcomings. Given the high complexity of analysing this game purely theoretically (as evidenced by the fact that no theoretical equilibrium results exist even for single double auctions), we analyse this problem by combining theory with a computational learning approach based on fictitious play to derive pure Bayes-Nash equilibria for traders' strategies in single and competing double auction marketplaces. We

[1] University of Southampton, UK, email: {bs07r,eg,pv,nrj@ecs.soton.ac.uk}

use this approach because it has been effectively applied to find equilibria in complex auction settings where traders' types are continuous [11]. We note that this is the first work that computes equilibria of both selection and bidding strategies for competing double auction marketplaces. In doing so, this paper advances the state of the art in the following ways. First, we formulate equations to calculate the trader's expected utility in our setting. Then we analyse the equilibrium bidding strategies of traders in a single double auction marketplace. We show empirically that the fictitious play algorithm converges to a unique pure Bayes-Nash equilibrium. We then analyse how market fees affect these strategies. By so doing, this is the first work that derives the equilibrium bidding strategies of traders in double auctions and analyses the effect of market fees on these strategies. Building on this analysis, we go on to study competing marketplace environments where we consider both the traders' bidding and market selection strategies. In particular, we analyse the traders' equilibrium strategies in two cases. In the first, traders are restricted to trade in one marketplace at a time, and we find that all traders who choose a marketplace eventually converge to the same one. We then consider a case which often exists in practice. Specifically, we allow one side (e.g. buyers) to choose multiple marketplaces, whereas the other side (e.g. sellers) can only choose one marketplace at a time. For this, we find that sellers converge to different marketplaces in equilibrium and they have more market power than buyers.

The structure of the paper is as follows. In Section 2, we introduce the basic setting of competing double auction marketplaces and derive the equations to calculate the trader's expected utility. In Section 3, we describe how to use fictitious play to derive the equilibrium of traders' strategies. Then, in Section 4, we analyse the traders' equilibrium strategies in detail. Finally, we conclude in Section 5.

2 Competing Double Auction Marketplaces

We first describe our basic setting for competing double auction marketplaces, before deriving equations for the trader's expected utility.

2.1 Basic Settings

We assume that there is a set of buyers, $\mathcal{B} = \{1, 2, ...B\}$, and a set of sellers, $\mathcal{S} = \{1, 2, ...S\}$. Each buyer and seller has a type[2], which is denoted as θ^b and θ^s respectively. We assume that the types of all buyers are independently drawn from the same cumulative distribution function F^b, with support $[0, 1]$, and the types of all sellers are independently drawn from the cumulative distribution function F^s, with support $[0, 1]$. The distributions F^b and F^s are assumed to be common knowledge and differentiable. The probability density functions are f^b and f^s respectively. In our setting, the type of each specific trader is not known to the other traders of the marketplaces, but the type distribution functions are common knowledge. In addition, we assume that there is a set of competing marketplaces $\mathcal{M} = \{1, 2, ...M\}$, that offer places for trade and provide a matching service between the buyers and sellers.

In our setting, we assume that marketplaces match buyers with sellers according to the *equilibrium matching* policy. We consider such a policy since it aims to maximise traders' profits and thus maximises the allocative efficiency for the marketplace. In detail, this policy will match the buyer with v-th highest offer with the seller with v-th lowest offer if the seller's offer is not greater than the buyer's offer. Since we consider marketplaces to be commercial enterprises that seek to make a profit, we assume they charge fees for their service as match makers. The fee structure of a marketplace m is defined as $\mathcal{P}_m = (p_m, q_m)$, $p_m \geq 0$ and $q_m \in [0, 1]$, where p_m is a registration fee charged to traders when they enter the marketplace, and q_m is a percentage fee charged on profits made by buyers and sellers (in the following, we refer to such fees as profit fees)[3]. Moreover, we assume that traders will incur a small cost ϵ when they choose any marketplace, so that they slightly prefer choosing no marketplace than choosing a marketplace and making no transaction (even if $p_m = 0$). Furthermore, the transaction price of a successful transaction in marketplace m is determined by a parameter $k_m \in [0, 1]$ (i.e. a *discriminatory k-pricing* policy), which sets the transaction price of a matched buyer and seller at the point determined by k_m in the interval between their offers.

Now we describe how a trading round proceeds. First, all marketplaces publish their fees and pricing parameters. Second, based on the observed fees and pricing parameters, traders select marketplaces. Third, traders submit their offers. Finally, after all traders have submitted their offers, the marketplaces match buyers and sellers according to their matching policy and then execute the transactions. The traders' behaviour of bidding and market selection is determined by their strategies, which map the set of types to a set of actions. We can see that the way in which traders select marketplaces and bid is important since this significantly affects market share and market profits of the competing marketplaces. Given this, in the following, we present the traders' behaviour of bidding and market selection in more detail.

Specifically, we call the offers of the buyers *bids* and the offers of the sellers *asks*. We assume that bids and asks are *discrete* and the sizes of allowable bids and allowable asks are finite[4]. The ranges of possible bids and asks constitute the *bid space* and *ask space* respectively. For convenience, we further assume that buyers and sellers have the same offer space, which is given by $\Gamma = \{0, \frac{1}{D}, \frac{2}{D}, ..., \frac{D-1}{D}, 1\} \cup \{\ominus\}$, i.e. the bid(ask) space comprises $D+1$ allowable bids(asks) and the action $\ominus$ means not choosing the marketplace. Given this, we can now define a trader's action in terms of bidding and market selection across multiple competing marketplaces. Specifically, we use $\delta = (d_1, d_2, ..., d_M)$ to represent an action, where the trader chooses marketplace m and offers d_m in this marketplace if $d_m \neq \ominus$. The set of all possible actions constitutes the *action space*. We consider two types of trading settings: multi-home and single-home. In the former, traders can participate in all marketplaces simultaneously and the action space is given by $\Delta_{multi} = \Gamma^M$. In the latter, traders are restricted to participate in one marketplace and the action space is given by $\Delta_{single} = \{\delta \in \Delta_{multi} : \exists i \in \mathcal{M}, \forall j \neq i, d_i \in \Gamma \; and \; d_j = \ominus\}$. In what follows, we use Δ as the general term for the action space.

2.2 The Trader's Expected Utility

After defining the traders' bidding and market selection actions, we now derive the equations to calculate a trader's expected utility given its type and its own action, and given the probability distribution of actions played by other traders. Specifically, we use $\Omega^b = \left(\omega_1^b, \omega_2^b, ..., \omega_{|\Delta|}^b\right)$, $\sum_{i=1}^{|\Delta|} \omega_i^b = 1$, to represent the probability distribution of (other) buyers' actions, where the probability of the action δ_i^b used by a buyer is ω_i^b, and $\Omega^s = (\omega_1^s, \omega_2^s, ..., \omega_{|\Delta|}^s)$ for the sellers' actions. Furthermore, we use $\delta^b = (d_1^b, d_2^b, ..., d_M^b)$ and $\delta^s = (d_1^s, d_2^s, ..., d_M^s)$ to denote a buyer's action and a seller's action respectively.

[2] The type of a buyer is its *limit price*, the highest price it is willing to buy the item for, and the type of a seller is its *cost price*, the lowest price it is willing to sell the item for.

[3] These two types of fees are common in the literature [1, 7]. Other types of fees, such as information fees, transaction fees, have similar effects, and can easily be incorporated in our setting.

[4] The traders' actions are always discrete in practice since the numeraire is discrete.

In what follows, we derive the expected utility of a buyer, but the seller's is calculated analogously. Specifically, we calculate the expected utility of a buyer with type θ^b adopting the action $\delta^b = (d_1^b, d_2^b, ..., d_M^b)$ given the other buyers' action distribution Ω^b and the sellers' action distribution Ω^s. First, we introduce four support functions[5]: $h_m^b(d)$ denotes the probability that the buyers' bids are strictly higher than d in marketplace m; $l_m^b(d)$ denotes the probability that the buyers' bids are strictly less than d in marketplace m (when buyers do not select this marketplace, this is equivalent to bidding less than d); similarly, $h_m^s(d)$ denotes the probability that the sellers' asks are strictly less than d in marketplace m and $l_m^s(d)$ denotes the probability that the sellers' asks are strictly higher than d in marketplace m (when sellers do not select this marketplace, this is equivalent to asking higher than d).

In our setting, we assume that, if the trader participates in multiple marketplaces, it can trade multiple items, i.e. its expected utility in these marketplaces is *additive*[6]. Then the buyer's expected utility over all marketplaces is the sum of its expected utility in each marketplace, which is given by:

$$\tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{m=1}^{M} \tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s)$$

where $\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s)$ is the expected utility of the buyer in marketplace m when it bids d_m^b, which is given by:

$$\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s) =$$
$$\begin{cases} 0 & \textbf{\textit{if}}\ d_m^b = \ominus \\ \sum_{z=0}^{B-1} \sum_{x=0}^{B-1-z} \rho^b(z, x, d_m^b) * \tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s | z, x) - p_m - \epsilon & \textbf{\textit{if}}\ d_m^b \neq \ominus \end{cases}$$

where

$$\rho^b(z, x, d_m^b) \quad = \quad \binom{B-1}{z, x} * h_m^b(d_m^b)^z$$
$$* \big(1 - h_m^b(d_m^b) - l_m^b(d_m^b)\big)^x * l_m^b(d_m^b)^{B-1-z-x}$$

is the probability that there are exactly z bids strictly higher than d_m^b, exactly x bids tying with d_m^b (excluding the bid itself), and exactly $B - 1 - z - x$ bids strictly less than d_m^b, and $\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s | z, x)$ is the buyer's expected utility when this occurs. Recall that p_m is the registration fee charged to the buyer when it enters the marketplace m and ϵ is the constant cost. We now calculate $\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s | z, x)$, which is given by:

$$\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s | z, x) = \frac{1}{x+1} * \sum_{v=z+1}^{z+x+1} \tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v)$$

Note that, at this moment, the position of the bid among all bids in marketplace m, v, can be anywhere between $z+1$ and $z+x+1$. We use a tie-breaking rule where each of these possible positions occurs with equal probability $1/(x+1)$. Now, $\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v)$ is the buyer's expected utility when its bid is v-th highest among all bids in marketplace m, which is given by:

$$\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v) = \sum_{j=1}^{|\Delta|} \tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s)$$

where $\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s)$ is the buyer's expected utility when it attempts to be matched with the ask d_{jm}^s of action $\delta_j^s = (d_{j1}^s, d_{j2}^s, ..., d_{jM}^s)$, which is given by:

$$\tilde{U}_m(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s) =$$
$$\begin{cases} 0 & \textbf{\textit{if}}\ d_m^b \le d_{jm}^s \ \textit{or}\ d_{jm}^s = \ominus \\ \sum_{r=0}^{v-1} \sum_{t=v-r}^{S-r} \rho^s(r, t, d_{jm}^s) * \tilde{U}_{m|j}(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s) & \textbf{\textit{if}}\ d_m^b > d_{jm}^s \end{cases} \quad (1)$$

where

$$\rho^s(r, t, d_{jm}^s) \quad = \quad \binom{S}{r, t} * h_m^s(d_{jm}^s)^r$$
$$* \big(1 - h_m^s(d_{jm}^s) - l_m^s(d_{jm}^s)\big)^t * l_m^s(d_{jm}^s)^{S-r-t}$$

is the probability that there are exactly r asks strictly less than d_{jm}^s and exactly t asks equal to d_{jm}^s (including the ask itself). Note that t should be at least equal to $v - r$ in Equation 1, thus $\sum_{r=0}^{v-1} \sum_{t=v-r}^{S-r} \rho^s(r, t, d_{jm}^s)$ actually gives the overall probability that the bid d_m^b is matched with the ask d_{jm}^s. Finally, $\tilde{U}_{m|j}(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s)$ is the buyer's expected utility when it is matched with the ask d_{jm}^s. This is given by:

$$\tilde{U}_{m|j}(\theta^b, \delta^b, \Omega^b, \Omega^s, v, \delta_j^s) = \theta^b - d_m^b + (d_m^b - d_{jm}^s) * k_m * (1 - q_m)$$

where $(d_m^b - d_{jm}^s) * k_m$ is the buyer's share of the observed trading surplus, which is the difference of the matched bid and ask, and q_m is the profit fee charged to traders.

3 The Fictitious Play Algorithm

In this section we describe how we can use fictitious play (FP) to find the equilibrium market selection and bidding strategies of traders in our setting. In the standard FP algorithm [10], opponents are assumed to play a fixed mixed strategy. Then by observing relative appearance frequencies of different actions, the player can estimate their opponents' mixed strategies, and take a best response. The observed frequencies of opponents' actions are termed *FP beliefs*. In each round, all players estimate their opponents' mixed strategies and update their FP beliefs, and play a best response to their FP beliefs. All players continually iterate this process until it converges. This algorithm has two types of convergence. First, it may converge to the pure strategy, which means that after a number of iterations, the best response strategy of each player is stable. At this moment, all players' best response strategies constitute a pure Nash equilibrium. Second, it may converge in FP beliefs. At this moment, the converged FP beliefs constitute a mixed Nash equilibrium.

However, the standard FP algorithm is not suitable for analysing Bayesian games in which there is incomplete information (i.e. where the player's type is not known to the other players). In such games, a strategy is a function mapping a player's type to an action. In the standard FP algorithm, by observing the frequency of opponents' actions, we cannot know the actual strategy of a player since we do not know which type performs which action. To ameliorate this, in [11], a generalised fictitious play algorithm was proposed to analyse games with continuous types and incomplete information. Using this algorithm, if the players' action space is finite, when the FP beliefs converge, they converge to a pure Bayes-Nash equilibrium[7]. Moreover, in such settings, it is known that a pure Nash equilibrium always exists. However, in [11], researchers only showed how to use

[5] These are calculated by taking the sum of the probabilities of actions whose corresponding offers in marketplace m satisfy the conditions defined by these functions.
[6] In the future, we will look at the cases with substitutable and complementary items.

[7] Or if they converge to beliefs, the equilibrium can be purified producing a pure Bayes-Nash equilibrium [9].

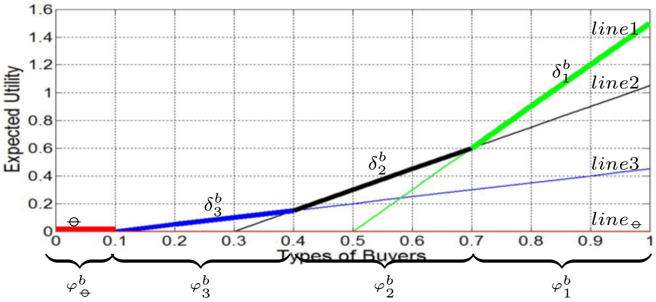

Figure 1. Piecewise linear expected utility functions.

this algorithm to analyse traders' strategies in single-sided auctions. Then in the following, given equations derived in Section 2.2, we apply this algorithm to find traders' equilibrium strategies in the much more complex competing double auction marketplaces environment.

Previously, we used Ω^b and Ω^s to denote the probability distributions of buyers' and sellers' actions respectively. In the FP algorithm, we use them to represent FP beliefs about the buyers'and sellers' actions respectively. Then given their beliefs, we need to compute the buyers' and the sellers' best response functions. In the following, we describe how to compute the buyers' best response function σ^{b*}, where $\sigma^{b*}(\theta^b) = argmax_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$. Considering the equations to calculate the buyer's expected utility in Section 2.2, we note that the buyer's expected utility $\tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$ is linear in its type θ^b for a given action. Now, the optimal utility that a buyer with type θ^b can achieve is $\tilde{U}^*(\theta^b) = max_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$. Given that the expected utility for each action is linear, and given a finite number of actions, the optimal function is the upper envelope of a finite set of linear functions, and thus is piecewise linear. An example with 4 actions, δ_1^b, δ_2^b, δ_3^b and $\ominus$, is given in Figure 1. Given each action, the buyer's expected utility with respect to its type is shown by $line1$, $line2$, $line3$ and $line_\ominus$ (i.e. x-axis) respectively. The optimal utility achieved by the buyer is represented by the set of thick piecewise linear segments. Each line segment corresponds to a type interval, where the best response action of each type in this interval is the same. In this figure, the best response action δ_i^b corresponds to the interval φ_i^b $(i = 1, 2, 3)$ and the best response action $\ominus$ corresponds to $\varphi_\ominus^b$. More generally, we can create the set of distinct intervals I^b, which constitute the continuous type space of buyers, i.e. $\bigcup_{\varphi^b \in I^b} \varphi^b = [0, 1]$, which satisfy the following conditions:

- For any interval φ^b, if $\theta_1^b, \theta_2^b \in \varphi^b$, then $\sigma^{b*}(\theta_1^b) = \sigma^{b*}(\theta_2^b)$, i.e. types in the same interval have the same best response action.
- For any distinct $\varphi_1^b, \varphi_2^b \in I^b$, if $\theta_1^b \in \varphi_1^b$, $\theta_2^b \in \varphi_2^b$, then $\sigma^{b*}(\theta_1^b) \neq \sigma^{b*}(\theta_2^b)$

Now we have computed the best response function and also provided the set of intervals of types corresponding to the best response actions. Based on this, we can calculate the inherent probability distribution of buyers' actions, which is down as follows. We know that given the buyers' type distribution function F^b and probability density function f^b, the probability that the buyer has the type in the interval φ^b is $\int_{\varphi^b} f(x)dx$, denoted by $F^b(\varphi^b)$. When the best response action corresponding to the interval φ_i^b is δ_i^{b*}, the probability that the action δ_i^{b*} is used by buyers is $\omega_i^b = F^b(\varphi_i^b)$. By calculating the probability of each action being used, we obtain the inherent action distribution of buyers, and we can then update the FP beliefs of buyers' actions, which is given by:

$$\Omega_{(\tau+1)}^b = \frac{\tau}{\tau+1} * \Omega_\tau^b + \frac{1}{\tau+1} * \Omega^b$$

where $\Omega_{(\tau+1)}^b$ is the updated FP beliefs of the buyers' actions for the next iteration round $\tau+1$, Ω_τ^b is the FP beliefs on the current iteration round τ, and Ω^b is the inherent action distribution of this round, which is absorbed by $\Omega_{(\tau+1)}^b$ with the standard update rate $\frac{1}{\tau+1}$. The computation of the sellers' best response function and belief updates is analogous. In our setting, we need to update both buyers' and sellers' FP beliefs simultaneously.

After introducing how to compute the best response function and update FP beliefs, we now give the structure of the FP algorithm (see Figure 2). In this algorithm, we measure the convergence in beliefs by calculating the convergence error, which is given by:

$$CE = max\left(max_{\delta_i^b \in \Delta}|\omega_{i(\tau+1)}^b - \omega_{i(\tau)}^b|, max_{\delta_i^s \in \Delta}|\omega_{i(\tau+1)}^s - \omega_{i(\tau)}^s|\right) * (\tau+1)$$

If $CE < \frac{1}{\tau}$, then the algorithm converges in beliefs, and a pure Bayes-Nash equilibrium of traders' strategies is reached.

Initial:
set iteration count $\tau = 0$
set the initial beliefs Ω_0^b and Ω_0^s
1. **loop**
2. Compute best response functions: $\sigma^{b*}(\theta^b)$ and $\sigma^{s*}(\theta^s)$
Generate the interval φ_i^b corresponding to the best response action δ_i^{b*}
Generate the interval φ_i^s corresponding to the best response action δ_i^{s*}
3. Compute inherent action distribution of buyers and sellers:
$\Omega^b = \{\omega_i^b | \omega_i^b = F^b(\varphi_i^b), i = 1, 2, ..., |\Delta|\}$
$\Omega^s = \{\omega_i^s | \omega_i^s = F^s(\varphi_i^s), i = 1, 2, ..., |\Delta|\}$
4. Update beliefs:
$\Omega_{(\tau+1)}^b = \frac{\tau}{\tau+1} * \Omega_\tau^b + \frac{1}{\tau+1} * \Omega^b$
$\Omega_{(\tau+1)}^s = \frac{\tau}{\tau+1} * \Omega_\tau^s + \frac{1}{\tau+1} * \Omega^s$
5. **if**(Convergence precision reached), **then**
6. **return** $\Omega_{(\tau+1)}^b$ and $\Omega_{(\tau+1)}^s$
7. **end if**
8. Set $\tau = \tau + 1$
9. **end loop**

Figure 2. The fictitious play algorithm.

4 Equilibrium Analysis of Traders' Strategies

In this section, we will use the FP algorithm to analyse the traders' equilibrium strategies. We first do so in a single marketplace. We then analyse the traders' bidding and market selection strategies in a competing marketplaces setting. In the following, for illustrative purposes, we show our results in a specific setting with 5 buyers and 5 sellers, and 11 allowable bids(asks) unless mentioned otherwise.

4.1 For a Single Marketplace

Many bidding strategies for double auctions have been proposed in the literature (such as AA [8], ZIP [3]). However, they all fail to answer what exactly traders should bid in equilibrium. This is important since how traders bid in a given marketplace will affect their expected utilities and this in turn their selection of marketplaces. Thus we first analyse a trader's bidding behaviour in a single marketplace.

We first consider a setting with no fees and set the small cost $\epsilon = 0.0001$. We also assume that $k_m = 0.5$, i.e. the transaction price is set in the middle of the matched bid and ask. Furthermore, we assume that buyers' types are independently drawn from a uniform function, and the same for sellers. Then we use the FP algorithm to analyse the traders' equilibrium strategies, and find that starting from different initial beliefs of traders' actions, all traders who choose the marketplace eventually converge to the same pure Nash equilibrium bidding strategy, which is shown in Figure 3. The light gray

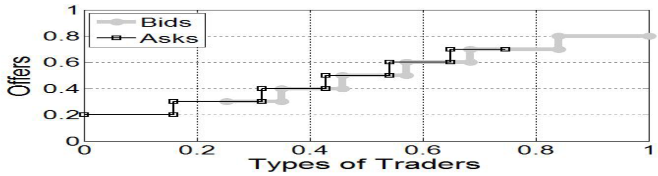

Figure 3. Equilibrium strategies with the same number of buyers and sellers.

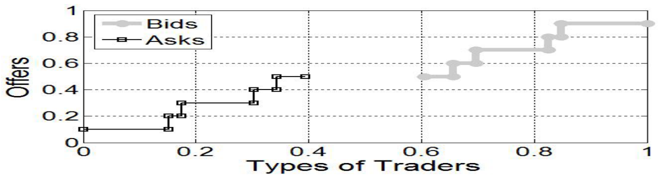

Figure 4. Equilibrium strategies of traders with registration fees.

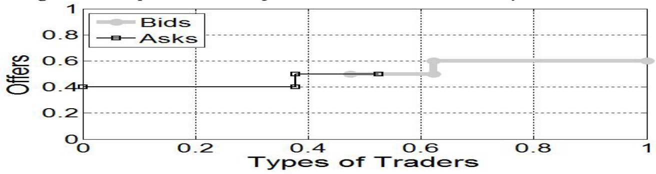

Figure 5. Equilibrium strategies of traders with profit fees.

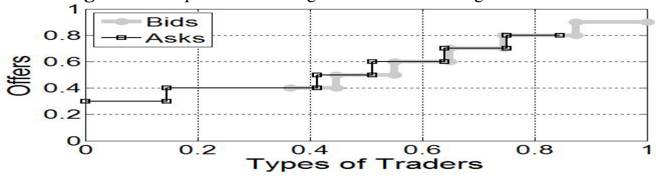

Figure 6. Equilibrium strategies with different numbers of buyers and sellers.

line with solid circles represents buyers' bids in equilibrium and the dark black line with squares represents sellers' asks in equilibrium. From this figure, we can see what traders will bid corresponding to their types in equilibrium. We also find that when buyers' types are lower than a certain point and sellers' types are higher than a certain point, they will not enter the marketplace because of the constant cost. These traders are called *extra-marginal* traders or *poor* traders since they cannot make any transactions. We also see that buyers with high types and sellers with low types (which we call *intra-marginal* traders or *rich* traders) shade their offers by bidding less or asking more in order to extract more profit from transactions.

Now we consider how fees can affect the traders' equilibrium bidding strategies. First we consider that the marketplace charges a registration fee. For example, if we assume that it charges a 0.1 registration fee, then the traders' equilibrium bidding strategies are shown in Figure 4. We can see that, compared to the case where no fees are charged (see Figure 3), there exists a bigger range of types of traders not choosing this marketplace, even including some buyers(sellers) whose types are higher(lower) than the equilibrium price of 0.5. This is because the registration fee causes negative profits for them. In addition, we further find that when registration fees are charged, rich buyers prefer to increase their bids, and thus increase the probability of being matched. In contrast, the buyers, whose types are relatively close to the equilibrium price 0.5, prefer to lower their bids in order to keep more profits.

Then we consider the case where the marketplace charges a profit fee. As an example, we assume that the marketplace charges a 50% profit fee. The results are shown in Figure 5. We still find that compared to the case where no fees are charged (see Figure 3), there exists a bigger range of types of traders not choosing the marketplace. However, compared to the case where registration fees are charged (see Figure 4), we can see that traders whose types are close to the equilibrium price still choose the marketplace. This is because the profit fee is a percentage fee charged on the observed trading surplus (which is the difference of the matched bid and ask), and will not cause negative profits for traders. Furthermore, we find that charging profit fees drives traders to bid close to the equilibrium price. This is because when profit fees are charged, marketplaces extract profits from traders according to their observed trading surplus. Then in order to keep more profits, the traders try to shade their offers by lowering bids or increasing asks to reduce the trading surplus observed by the marketplace.

In the above, we considered the case with the same number of buyers and sellers. We now analyse what happens to the traders' bidding strategies when there are different numbers of buyers and sellers. For example, we consider the case with 8 buyers and 5 sellers. Then the traders' equilibrium bidding strategies are shown in Figure 6. Compared to Figure 3, as can be expected, we see that because there are more buyers than sellers, the competition between buyers is more severe, and thus they have to raise their bids. For sellers, since they have a higher probability of being matched, they raise their asks. At this moment, sellers have more market power, and can therefore extract more profit from their transactions.

4.2 For Competing Marketplaces

In the above, we analysed how traders will bid in a single marketplace in equilibrium. Now we analyse how traders select marketplaces and bid in the competing marketplaces environment. As mentioned previously, there are two types of trading settings: single-home and multi-home. Furthermore, we also consider a hybrid setting where one side can only participate in one marketplace, where the other side can participate in multiple marketplaces. Specifically, we consider the case where buyers can participate in multiples marketplaces and sellers can only participate in one marketplace at the same time. For example, a seller wants to sell one item through online auctions. Once it chooses Amazon or Ebay, he is committed to sell the item in this specific auction. He cannot choose other online auctions at the same time. However, for buyers, they can place bids in multiple online auctions at the same time to find the best deal. Note that the results of the opposite case where sellers can choose multiple marketplaces and buyers can only choose one are identical since the marketplaces are symmetric.

4.2.1 Single-Home Trading

We now consider traders' equilibrium strategies when there are two competing marketplaces[8]. Other settings are the same as above. We first consider the case where the marketplaces charge no fees to traders. By using FP, we find that, except for the extra-marginal traders choosing no marketplace, all other traders eventually converge to one marketplace in equilibrium. Since the two marketplaces are identical at this moment, the traders will eventually converge to marketplace 1 or 2 with the same probability. In addition, we find that the traders' bidding strategies in the converged marketplace are the same as the case with a single marketplace (i.e. Figure 3).

We now consider what happens to the traders' behaviour when they are charged fees. First we consider the cases where both marketplaces charge registration fees or both charge profit fees. We run simulations with many possible initial beliefs, and find that traders eventually converge to one marketplace, which depends on initial FP beliefs and market fees. Since traders converge to one marketplace, the equilibrium bidding strategies are the same as the case with a single marketplace (i.e. Section 4.1).

[8] We also considered the case with more than two competing marketplaces. However, the results are similar to the case with two competing marketplaces and therefore we omit them.

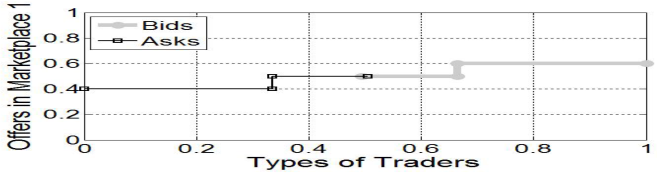

Figure 7. Equilibrium strategies of traders with different types of fees.

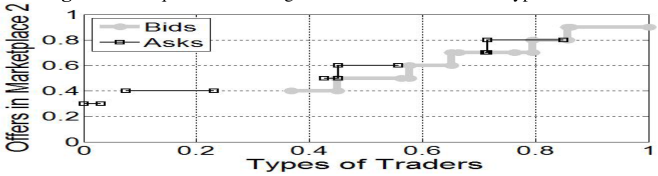

Figure 9. Equilibrium strategies of traders in marketplace 2 in the hybrid setting.

Now we consider the case where marketplace 1 charges a profit fee and marketplace 2 charges a registration fee. For example, we consider the profit and registration fees with a little extreme, where marketplace 1 charges a very high profit fee of 0.9, and marketplace 1 charges a registration fee of 0.1. If initial beliefs are uniform (i.e. all actions are equally probable), we find that all traders eventually converge to marketplace 1, and the equilibrium bidding strategies are shown in Figure 7. The reason of traders converging to marketplace 1, which seems more expensive, is as follows. When a high profit fee is charged, the traders shade their offers more to keep profits (as can be seen in Figure 7). However, shading has no effect in the case of registration fees. Therefore, traders will prefer the marketplace charging profit fees compared to registration fees.

We also run simulations with many other fee combinations. We always find that all traders converge to one marketplace. This is because when traders split into different marketplaces, a trader's probability of being matched decreases, which results in the loss of the trader's profit. Thus traders want to concentrate in one marketplace.

4.2.2 *Multi-Home and Hybrid Trading*

In the above, we analysed traders' strategies when they can only select one marketplace at a time. When multi-home trading is allowed, the competition between marketplaces is reduced. This is because traders will choose both marketplaces if their expected profits in each marketplace are positive. Since we assumed that the expected utilities of different marketplaces are additive (Section 2.2), all traders, except for the extra-marginal traders choosing no marketplace, will choose both marketplaces. The analysis of their bidding strategies is then identical to Section 4.1.

Now we consider a hybrid case where buyers can participate in multiple marketplaces and sellers can only choose one marketplace at a time. Considering the same setting as before and without fees, the results are shown in Figures 8 and 9. From these figures, we can see that sellers eventually split and place asks in different marketplaces in equilibrium. Now two competing marketplaces co-exist. This co-existence is caused by internal competition between sellers. In the marketplace, the sellers have to compete with each other in order to be matched with buyers and make transactions, and thus they prefer those marketplaces with fewer sellers. Because identical buyers stay in both competing marketplaces in this case, then the attractiveness from the buyers to the sellers in both marketplaces is the same. At this moment, the internal competition between the sellers takes effect, which drives the sellers to stay in different marketplaces. Another interesting phenomenon is that compared to the traders' equilibrium strategies in Figure 3, we find that buyers raise their bids and sellers also raise their asks in this case. The reason is as follows. As the sellers are split in two marketplaces, then in each marketplace the

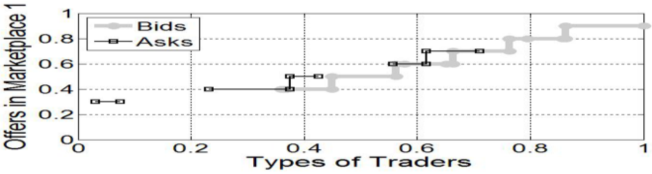

Figure 8. Equilibrium strategies of traders in marketplace 1 in the hybrid setting.

number of sellers is less than the number of buyers. Thus as per our previous analysis (see Figure 6), sellers have more market power than buyers, and so buyers raise their bids in order to be matched and sellers raise their asks to extract more profits from transactions.

5 Conclusions

In this paper, we used a FP algorithm to analyse the equilibrium strategies of traders with continuous types in competing double auction marketplaces. Specifically, we first analysed traders' equilibrium bidding strategies in a single marketplace, where we found that traders shade their offers in equilibrium. We further analysed the effect of market fees on traders' equilibrium bidding strategies. Then in the competing marketplaces environment, we analysed the traders' equilibrium strategies for market selection and bidding in two cases. In the first, where traders can participate in one marketplace at a time, we showed that traders eventually converge to one marketplace in equilibrium. In the second case where buyers can participate in multiple marketplaces and sellers can only choose one marketplace at the same time, we found that sellers will eventually converge to different marketplaces. We also showed that in this case, sellers have more market power than buyers.

In this paper, we have analysed how market fees can affect traders' equilibrium strategies. In the future, we intend to analyse how competing marketplaces should set their fees to make profits. Specifically, because traders shade their offers and will shade more when a profit fee is charged, then the marketplace may not be able to obtain the target profit even though it charges a very high profit fee. On the other hand, charging a high registration fee may cause negative profits for lots of traders, and thus drive them to leave the marketplace. Given this, we would like to analyse what types of fees, or what combinations of different types of fees, are effective at allowing the marketplace to obtain the target profit.

REFERENCES

[1] Caillaud B. and Jullien B., 'Chicken & egg: Competition among intermediation service providers', *The RAND Journal of Economics*, **34(2)**, 309–328, (2003).
[2] Shi B., Gerding E. H., Vytelingum P., and Jennings N. R., *A game-theoretic analysis of market selection strategies for competing double auction marketplaces*, 857–864, 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS2010), 2010.
[3] Cliff D. and Bruten J., 'Minimal-intelligence agents for bargaining behaviors in market-based environments', Technical report, Hewlett Packard Research Laboratories, (1997).
[4] Friedman D. and Rust J., *The double auction market: institutions, theories and evidence*, volume XIV of Santa Fe Institute Studies in the Science of Complexity, Perseus Publishing, 1993.
[5] Gerding E. H., McBurney P., Niu J., Parsons S., and Phelps S., 'Overview of CAT: A market design competition', Technical report, Dept. of Computer Science, University of Liverpool, (2007).
[6] Rochet J. and Tirole J., 'Platform competition in two-sided markets', *Journal of European Economic Association*, **1(4)**, 990–1029, (2003).
[7] Armstrong M., 'Competition in two-sided markets', *RAND Journal of Economics*, **37**, 668–691, (2006).
[8] Vytelingum P., Cliff D., and Jennings N. R., 'Strategic bidding in continuous double auctions', *Artificial Intelligence Journal*, **172(14)**, 1700–1729, (2008).
[9] Radner R. and Rosenthal R. W., 'Private information and pure-strategy equilibria', *Mathematics of Operations Research*, **7(3)**, 401–409, (1982).
[10] Brown G. W., 'Iterative solutions of games by fictitious play', *Activity Analysis of Production and Allocation*, (1951).
[11] Rabinovich Z., Gerding E. H., Polukarov M., and Jennings N. R., *Generalised fictitious play for a continuum of anonymous players*, 245–250, 21st International Joint Conference on Artificial Intelligence (IJCAI), 2009.