

An Equilibrium Analysis of Trading Across Multiple Double Auction Marketplaces using Fictitious Play

Bing Shi^a, Enrico H. Gerding^b, Nicholas R. Jennings^{c,d}

^a*School of Computer Science and Technology, Wuhan University of Technology, Wuhan, 430070, China*

^b*School of Electronics and Computer Science, University of Southampton, Southampton, SO17 1BJ, UK*

^c*Department of Computing, Imperial College London, London, SW7 2AZ, UK*

^d*Department of Computing and Information Technology, King Abdulaziz University, Saudi Arabia*

Abstract

We investigate how automated traders strategically select marketplaces and submit offers across multiple double auction marketplaces. We model the problem as a Bayesian game with traders that have continuous private values, and use fictitious play to analyse the traders' Nash equilibrium market selection and bidding strategies. We do this for different trading environments (*isolated*, *single-home*, *multi-home* and *hybrid*) and different types of goods (*independent*, *substitutable* and *complementary*). We find that, in an isolated marketplace, the fictitious play algorithm converges to a Bayes-Nash equilibrium. In the single-home setting, all traders eventually converge to the same marketplace and the setting reduces to that of an isolated marketplace. In the multi-home setting with perfectly substitutable goods, buyers with high values only bid in one marketplace, whereas buyers with low values bid in multiple marketplaces. Then, for perfectly complementary goods, only buyers with high values bid in multiple marketplaces and buyers with low values enter no marketplaces. Finally, in the hybrid setting with perfectly complementary goods, traders choose no marketplaces.

Keywords: Trading Strategies, Double Auctions, Bayes-Nash Equilibrium, Fictitious Play

Email addresses: bingshi@whut.edu.cn (Bing Shi), eg@ecs.soton.ac.uk (Enrico H. Gerding), n.jennings@imperial.ac.uk (Nicholas R. Jennings)

1. Introduction

Electronic exchanges, in which securities, futures, stocks and commodities can be traded, are becoming ever more prevalent. In addition to various electronic financial markets (such as NASDAQ and NYSE), electronic betting exchanges have emerged which allow gamblers to trade online (such as World Bet Exchange <http://www.wbx.com> and BetFair <http://www.betfair.com>). Many of these adopt the double auction market mechanism which is a type of two-sided marketplace for matching buyers (one side) and sellers (the other side) (Friedman and Rust, 1993). The advantages of this type of mechanism are that traders can enter the marketplace at any time and they can trade multiple items in one place without travelling around several marketplaces. Furthermore, this mechanism is highly efficient in economic terms (Smith, 1962).

Now, often electronic exchanges do not exist in isolation, and there are several competing electronic exchanges where traders can participate. In finance, companies can be listed on multiple electronic exchanges. In addition, alternative automated trading systems, often called “dark pools” or “dark liquidity”, are propagating rapidly (Carrie, 2008). For online betting, gamblers can choose to trade bets in one or multiple electronic bet exchanges, such as World Bet Exchange, Betfair and Betdaq (<http://www.betdaq.com>). In these markets, software agents are often used to make autonomous trading decisions because of the speed of trading that is required (Ma and Leung, 2007; Vytelingum et al., 2008; Dang et al., 2015). However, as there are multiple such markets, these trading agents need to both select which ones to participate in, as well as how much to bid in each one selected.

Against this background, in this paper we use a computational learning approach called *fictitious play* to analyse how trading agents behave strategically in terms of *selecting marketplaces* and *submitting offers* in the context of multiple double auction marketplaces. Intuitively, trading agents’ strategies will be determined by their own preferences (types), which are usually heterogeneous and privately known. Moreover, traders’ strategies are affected by the trading environments. Specifically, we distinguish between four different trading environments. The first is the *isolated* marketplace, which corresponds to the commonly-studied setting with no competing marketplaces. The second is *single-home trading* where both buyers and sellers can only select one of the available marketplaces. The third is *multi-home trading* where both buyers and sellers can participate in multiple marketplaces simultaneously. Finally, in *hybrid trading* settings one side of traders can only enter one marketplace (i.e. single-home trading), while the other

side of traders can enter multiple marketplaces (i.e. multi-home trading).

It is known that different trading environments have different effects on the strategic behaviour of the traders. For example, in a single-home environment, traders will only participate in the most profitable marketplace. However, in the multi-home case, traders will participate in any marketplace that provides non-negative (or positive) profits. In addition to the impact of the trading environments, the types of the goods traded between buyers and sellers can also affect their behaviour. Specifically, when multiple goods are traded across multiple marketplaces, these goods can be either *independent*, *substitutable* or *complementary*. When they are independent, the trader's valuation for the multiple goods is additive, i.e. equal to the sum of its valuation for each individual good. When they are substitutable, the trader's valuation is subadditive, i.e. less than the sum of its valuation for each individual good. When the goods are complementary, the trader's valuation is superadditive, i.e. greater than the sum of its valuation for each individual good. These different types also affect traders' strategies. For example, when trading complementary goods, buyers may prefer to buy as many as they can, and thus will try to bid high in several marketplaces to maximise the number of transactions. In this paper, we will consider all of the above factors and derive traders' market selection and bidding strategies in different trading environments with different types of goods.

Intuitively, we can see that the decision about which marketplace to select and how much to offer depends on other traders' decisions. Therefore, game theory (Fudenberg and Tirole, 1991), which mathematically studies such strategic interactions between self-interested agents (where an individual's success in making choices depends on the choices of others), is the appropriate tool to analyse our system. Indeed, game theory has been widely used to analyse the strategic behaviour of traders in similar areas, see, for example, Lin and Chou (2004); Yuan and Zeng (2012) and Shi et al. (2013). Specifically, we assume that traders have a privately known value or *type* which determines their utility for the obtained goods, and this type is i.i.d. sampled from a continuous probability distribution. A trader does not know the types of other traders, but knows the probability distribution from which the types are sampled. Given this, the appropriate solution concept is the Bayes-Nash equilibrium in which each trading agent makes a best response against the other agents' strategies. Due to the high complexity of the game we study, it is very difficult, perhaps impossible, to derive the equilibrium solution by purely theoretical approaches based on lemmas or theorems. Therefore, we use numerical approaches to compute the equilibrium. Furthermore, this game involves a continuous trader type space, which results in an infinite game.

Therefore, standard approaches, such as the Lemke-Howson algorithm (Lemke and Howson, 1964) or the Govindan-Wilson algorithm (Govindan and Wilson, 2003), cannot be used. Instead, we adopt a computational learning approach based on fictitious play (FP) to approximate pure Bayes-Nash equilibria for the traders' strategies. We use this approach because it has previously been effectively applied to find equilibria in complex auction settings where traders' types are continuous (Rabinovich et al., 2013).

In so doing, this is the *first* work to analyse traders' Nash equilibrium market selection and bidding strategies across multiple double auction marketplaces that takes into account different trading environments with different types of goods. In more detail, the contributions of this paper are as follows. Firstly, we analyse the equilibrium bidding strategies in isolated double auction marketplaces. Despite the extensive existing research studying these auctions, computing equilibria for this setting is still challenging (although the existence of the Nash equilibrium has been shown and the solution for the simple setting with only one buyer and one seller has been analysed, see Section 2). We then go on to study the traders' market selection and bidding strategies across multiple double auction marketplaces. In single-home trading environments, we find that all traders eventually converge to the same marketplace. In the multi-home environment, when trading perfectly substitutable goods, buyers having high values will only bid in one marketplace, and buyers having low values will bid in multiple marketplaces to increase the probability of being matched. When trading perfectly complementary goods, only buyers that have high values will bid in multiple marketplaces, while buyers with low values will not enter any marketplace. Furthermore, in the hybrid trading environment with perfectly substitutable goods, all traders only choose one marketplace in equilibrium and, as buyers' values for multiple goods increase, buyers will increasingly bid in multiple marketplaces, which causes sellers to participate in different marketplaces depending on their types. For perfectly complementary goods, we find that buyers choose no marketplaces since they incur a high risk of loss, and so neither will sellers. However, as buyers' values for individual items increase or buyers have more market power than sellers, we find that some trade does occur.

The structure of the rest of the paper is as follows. In Section 2, we discuss related work. In Section 3, we describe the setting for analysing traders' strategies across multiple double auction marketplaces, and derive the expected utilities of traders in this setting. In Section 4, we describe the FP algorithm used in our analysis. In Section 5, we use this algorithm to analyse traders' equilibrium strategies. Finally, we conclude and discuss the limitations in Section 6.

2. Related Work

We start by discussing related work on analysing trading strategies across multiple single-sided auctions. We then introduce related work on bidding strategies in isolated double auction marketplaces. Finally, we describe work on analysing traders' behaviour in the context of multiple double auction marketplaces.

2.1. Trading Across Multiple Single-Sided Auctions

A number of works analyse trading agents' strategies across multiple single-sided auctions with independent goods, such as English auctions, Dutch auctions and sealed-bid auctions (Preist et al., 2001; Byde et al., 2002; Stone et al., 2003; Anthony and Jennings, 2003; He et al., 2006; Dang et al., 2015). In this context, in order to make effective trading decisions, buyers need to determine the best set of auctions in which to bid and determine how much to bid in the chosen auctions. Existing research mainly adopts empirical approaches to design buyers' bidding strategies across multiple single-sided auctions. Specifically, Anthony and Law (2012) analyse how sellers set reserve prices in multiple English auctions heuristically. Furthermore, Ashlagi et al. (2013) investigate the bidding strategies of advertisers across multiple ad auctions. In addition to analysing bidding across multiple auctions with independent goods, Zeng et al. (2004) analyse the bidding strategies on goods with combinatorial interdependencies in the setting of multiple posted-price markets and multiple sealed-bid auctions respectively. Furthermore, Gerding et al. (2008) analyse the optimal bidding strategy across multiple simultaneous Vickrey auctions with perfectly substitutable goods in a theoretical way. Wellman et al. (2008) and Goeree and Lien (2014) analyse bidding strategies across multiple ascending auctions by considering complementary and substitutable goods. Wellman et al. (2012) and Mayer et al. (2013) analyse bidding strategies across multiple sealed-bid auctions by considering complementary and substitutable goods. In contrast, instead of analysing trading across multiple single-sided auctions, we focus on the market selection and bidding strategies across multiple double auctions. In double auctions, transactions can happen between multiple buyers and multiple sellers. Thus, an agent's optimal bidding strategy not only depends on the bids of their direct competitors, but also on the agents in the other side of the market. Furthermore, our setting considers different trading environments with multiple double auction marketplaces, which also increases the complexity of the analysis.

2.2. Trading in an Isolated Double Auction

A number of heuristic bidding strategies have been proposed for traders bidding in an isolated double auction, such as Gode and Sunder (1993); Gjerstad and Dickhaut (1998); Ma and Leung (2007) and Vytelingum et al. (2008). Furthermore, Phelps et al. (2010) use evolutionary game theory to investigate the Nash equilibrium given a restricted strategy space. However, these restricted strategies do not necessarily constitute a Nash equilibrium when considering the entire space of possible strategies. Moreover, these strategies are not suitable for traders to bid across multiple marketplaces because they do not consider the choice of marketplaces.

In addition to these heuristic strategies, there also exists theoretical work on investigating the existence of equilibrium bidding strategies in isolated double auction marketplaces. In double auctions there exists a trivial equilibrium strategy where buyers bid zero and sellers ask for the maximal allowed offer. Since there is a zero probability of transactions happening between buyers and sellers in this trivial equilibrium, this equilibrium is meaningless. Therefore, researchers focus on non-trivial equilibria in which there is a positive probability of transactions happening. The first such work is by Chatterjee and Samuelson (1983). They show that, in the case of one buyer and one seller, an equilibrium solution exists where the offers of the buyer and the seller are monotonic in the respective reservation prices (types). A number of other works extend the above. Williams (1991) analyses the buyer's Nash equilibrium bidding strategy in a double auction where the seller has a dominant strategy of bidding its true value. By fixing the seller's strategy, the authors prove the existence of the buyer's Nash equilibrium bidding strategy. Then Jackson and Swinkels (2005) show the existence of a non-trivial mixed-strategy equilibrium for double auctions in a variety of settings given a finite set of offers, and Reny and Perry (2006) show that, when there are sufficiently many buyers and sellers with a finite set of discrete offers, there exists a monotonic pure equilibrium bidding strategy for traders. Although all of these works show the *existence* of non-trivial equilibrium bidding strategies in double auctions, they do not indicate exactly what traders will bid in equilibrium when there exist multiple buyers and sellers. Furthermore, there exists no theoretical work on analysing equilibrium bidding strategies across multiple double auction marketplaces. In addition to these works, Bredin and Parkes (2005) propose a framework for designing truthful double auction marketplaces, in which traders will bid truthfully (i.e. bid their types). However, the auctions that are used in practical settings are not truthful, and traders usually have an incentive to shade their bids in order to make more profits.

2.3. Trading Across Multiple Double Auctions

In the context of multiple double auctions, Cai et al. (2010) experimentally analyse how standard economic measures are affected by the presence of multiple marketplaces when traders select marketplaces and submit offers in a heuristic way, and then Miller and Niu (2012) experimentally analyse traders' market selection strategies in the single-home trading environment. In contrast, in our work, we analyse both the Nash equilibrium market selection and bidding strategies in different trading environments with different types of goods. Similar to this work, Shi et al. (2010; 2013) analyse the Nash equilibrium market selection strategy in the context of multiple double auction marketplaces. However, these works only consider a limited number of discrete trader types and assume that traders can only choose one marketplace at a time (i.e. single-home trading). Furthermore, these works do not solve the equilibrium bidding strategy. In contrast, we address these limitations by considering a game with continuous trader types and analyse both the equilibrium market selection and bidding strategies in different trading environments with different types of goods. This new setting with continuous trader types requires a fundamentally different approach to address the Nash equilibrium strategy and also enables us to considerably enrich the analysis.

Finally, game theory has often been used to analyse the strategic interactions in electronic marketplaces (Lin and Chou, 2004; Yuan and Zeng, 2012; Shi et al., 2013). In this paper, as we discussed previously, we will use game theory to analyse the traders' Nash equilibrium market selection and bidding strategies. Furthermore, due to the high complexity of the game analysed in this paper, we will use fictitious play to approximate the traders' equilibrium market selection and bidding strategies. However, the traditional fictitious play algorithm (von Neumann and Brown, 1950; Brown, 1951) 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) and where the types are continuous (resulting in an infinite game). To address this, Rabinovich et al. (2013) provided a generalised fictitious play algorithm to analyse Bayesian games with continuous types and a finite action space. However, they only show how to use this algorithm to analyse traders' strategies in single-sided auctions. We extend their fictitious play algorithm to analyse the equilibrium market selection and bidding strategies for traders in the much more complex environment of multiple double auction marketplaces.

3. The Double Auction Marketplace Framework

We first introduce the basic setting for analysing traders' strategies across multiple double auction marketplaces. We then proceed to the first step of the theoretical analysis and derive the equations to calculate the expected utilities of the traders in this setting. This will then be used by the FP algorithm in Section 4 to compute the best response action and approximate the Bayes-Nash equilibrium market selection and bidding strategy.

3.1. Basic Setting

We assume there is a set of buyers, $\mathcal{B} = \{1, 2, \dots, B\}$, and a set of sellers, $\mathcal{S} = \{1, 2, \dots, S\}$. Each buyer and each seller can only trade a single unit of the good in one marketplace. All goods are identical. Each buyer and seller has a type,¹ which is denoted as θ^b and θ^s respectively. We assume that the types of all buyers are i.i.d drawn from the cumulative distribution function F^b , with support $[0, 1]$, and the types of all sellers are i.i.d 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 private information, i.e. not known to the other traders.

In addition, we assume that there is a set of double auction marketplaces $\mathcal{M} = \{1, 2, \dots, M\}$, that offer places for trade and provide a matching service between the buyers and sellers. Traders will incur a small cost ι when they choose any marketplace (representively, for example, the time cost for trading online or travel and time costs for trading in shopping malls). We do this so that they slightly prefer choosing no marketplace to choosing a marketplace but not transacting.

Furthermore, we assume that marketplaces adopt a clearing house double auction market mechanism, which means that the marketplaces match buyers with sellers when all traders have submitted their offers. We also assume that marketplaces match buyers with sellers according to the *equilibrium matching* policy. 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. By adopting the clearing house mechanism and the equilibrium matching policy, the marketplaces can match traders in a highly efficient way. Moreover, the transaction price of a successful transaction in marketplace m is determined by a param-

¹The type of a buyer is its *limit price*, the highest price it is willing to buy the item for. The type of a seller is its *cost price*, the lowest price it is willing to sell the item for.

eter $k_m \in [0, 1]$ (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.

After describing the basic settings for traders and marketplaces, we now introduce the trader’s value or values for possessing multiple goods. As mentioned above, we assume that each trader can only trade one single unit of the good in each marketplace. However, when multi-home trading is available, traders may trade multiple goods when they participate in multiple marketplaces. These goods can be either independent, substitutable, or complementary.² We need to describe these different preferences, which are modeled as follows. For a buyer with type θ^b , the valuation that it derives when it successfully purchases T units of the good is given by:

$$v^b(\theta^b, T) = \alpha^b(T) \times \theta^b \quad (1)$$

where we refer to $\alpha^b(T)$ as the “buyer preference coefficient function”, determining whether the buyers have independent, substitutable or complementary preferences. In more detail, if the goods are independent for the buyer, then the total valuation for the individual goods is additive, i.e. $\alpha^b(T) = T$. If, on the other hand, the goods are substitutable, then the total valuation for getting T goods is subadditive, i.e. $\alpha^b(1) = 1$ and $0 \leq \alpha^b(T) - \alpha^b(T - 1) < 1, \forall T \geq 2$. In particular, when the goods are perfectly substitutable for the buyer, we have $\alpha^b(T) = 1, T \in \{1, \dots, M\}$. Finally, if goods are complementary, the total valuation is superadditive, i.e. $\alpha^b(1) = 1$ and $\alpha^b(T) - \alpha^b(T - 1) > 1, \forall T \geq 2$. Specifically, for perfectly complementary goods, we have $\alpha^b(T) = 0, T \in \{1, \dots, M - 1\}$ and $\alpha^b(M) = 1$, i.e. the buyer obtains valuation θ^b when it successfully purchases M goods, and the buyer obtains zero valuation when it purchases less than M . The valuation function for a seller is defined analogously, where $\alpha^s(T)$ is the seller preference coefficient function. We assume that $\alpha^b(T)$ and $\alpha^s(T)$ are the same for all buyers and sellers, and these parameters are common knowledge.

We now describe the actions that traders take in this setting (i.e. selecting marketplaces and submitting offers). We call the offers of the buyers *bids* and the offers of the sellers *asks*. Specifically, we make the assumption that there are a

²Note that in this paper we assume that all goods are identical. Although the notion of substitutable and complementary goods usually refers to different types of goods, in some situations identical goods can be substitutable or complementary. For example, when a couple purchases cinema tickets, the valuation of two tickets is intuitively higher than the sum of the valuation on each individual ticket (i.e. complementary). When a single person buys two cinema tickets, the valuation of two tickets for this person is equal to the valuation of one ticket (i.e. substitutable).

finite number of bids and asks and that these are discrete. The reason for doing so is two-fold. First, this assumption is more realistic than having continuous bids because, in practice, the numeraire is discrete. Second, it allows us to compute the Bayes-Nash equilibrium more easily using fictitious play. Now, the ranges of possible bids and asks constitute the *bid space* and *ask space* respectively. In reality, buyers and sellers usually have the same offer space. In our setting, the offer space for buyers and sellers is denoted by $\Psi = \{0, \frac{1}{D}, \frac{2}{D}, \dots, \frac{D-1}{D}, 1\} \cup \{\emptyset\}$, i.e. the offer space comprises $D + 1$ allowable offers from 0 to 1 with step size $1/D$, and \emptyset means not submitting an offer in the marketplace (i.e. not entering the marketplace). In our setting, traders may participate and place offers in multiple marketplaces. We refer to a combinational offer across multiple marketplaces as an *action*. Formally, a buyer's action is defined as a tuple $\delta^b = \langle d_1^b, d_2^b, \dots, d_M^b \rangle$, where the buyer bids d_m^b in marketplace m if $d_m^b \neq \emptyset$, and does not choose marketplace m if $d_m^b = \emptyset$. Note that this definition actually combines the buyer's market selection and the bids in the selected marketplaces as a whole. Similarly, a seller's action is given by $\delta^s = \langle d_1^s, d_2^s, \dots, d_M^s \rangle$. The set of all possible actions constitutes the *action space*, which is defined as $\Delta = \Psi^M$. Note that in our system, both buyers and sellers have the same action space.

Now, a trader's action choice depends on its type. Hence, a *strategy*, is defined as a mapping from the set of types to the action space. Formally, we use $\sigma^b : [0, 1] \rightarrow \Delta$ and $\sigma^s : [0, 1] \rightarrow \Delta$ to denote the buyer and the seller trading strategies respectively. Note that the expected utility of a trader is directly dependent on its beliefs about the other traders' action choices. Therefore, instead of looking at traders' strategies, in what follows, the expected utility is expressed directly in terms of the traders' action distributions. This will also be convenient when we use fictitious play to derive traders' equilibrium strategies (see Section 4). Specifically, we use ω_i^b to denote the probability of action δ_i^b being chosen by a buyer, and ω_i^s to denote the probability of action δ_i^s being chosen by a seller. Furthermore, we use $\Omega^b = (\omega_1^b, \omega_2^b, \dots, \omega_{|\Delta|}^b)$, $\sum_{i=1}^{|\Delta|} \omega_i^b = 1$, to represent the probability distribution of buyers' actions and $\Omega^s = (\omega_1^s, \omega_2^s, \dots, \omega_{|\Delta|}^s)$ for the sellers' action distribution. Note that, given a trader's strategy and the type distribution function, we can derive the probability of a certain action being played by the trader. Specifically, we use $\sigma^{-b}(\delta_i^b) \subseteq [0, 1]$ to denote the set of buyer types adopting action δ_i^b . Then the probability of the action δ_i^b being played is $\omega_i^b = \int_{x \in \sigma^{-b}(\delta_i^b)} f^b(x) dx$. The calculation of the seller action distribution is analogous.

3.2. The Trader's Expected Utility

Before analysing the strategies of the traders, we first need to derive the equations to calculate their expected utilities, which are defined as the expected profits that traders can make in the marketplaces. In what follows, we derive the expected utility of a buyer, but the seller's is calculated analogously. We can see that a buyer's expected utility depends on its type, its own action and its beliefs about the action choices of other traders. In the following, we calculate the expected utility of a buyer with type θ^b adopting the action $\delta^b = \langle d_1^b, d_2^b, \dots, d_M^b \rangle$ given the other buyers' action distribution Ω^b and the sellers' action distribution Ω^s . The expected utility consists of two parts: the *expected valuation* of the goods and the *expected payment* made for the goods. In the following, we derive these two parts separately.

3.2.1. The Expected Valuation

Firstly, we derive the buyer's expected valuation. In our setting, since we consider the equilibrium matching policy (where buyers submitting high bids are matched with sellers submitting low asks), a buyer's expected valuation is determined by its position in the available marketplaces (the buyer's position is the rank of the buyer's bid among all descendingly sorted bids) and the sellers' action choices. Furthermore, it is also determined by the number of units it can win. In the following, we calculate the expected position of the buyer when adopting action δ^b and then compute the probability of the specific numbers of sellers choosing different actions. Finally, given the buyer's position and sellers' action choices, we derive the buyer's expected valuation by considering all possible numbers of units it can win.

The Buyer's Expected Position. Clearly, a buyer's position in a marketplace is determined by its own action and those of other buyers. However, since we do not know the actions of other traders, we can only derive probabilistic information about the buyer's position in a marketplace. Furthermore, the bids placed in multiple marketplaces from the same action are correlated with each other. Thus when the buyer chooses a specific action, its positions in different marketplaces are also correlated with each other. Therefore, we cannot consider each marketplace independently, rather we need to consider the buyer's *joint positions* in marketplaces. In addition, since different buyers may place the same bids in the same marketplace, we need to have a tie-breaking rule to determine a buyer's position.

In the following, we describe how to determine a buyer's joint positions when it takes action δ^b . Intuitively, by knowing the number of buyers choosing each ac-

tion, we know what exact bids are placed in each marketplace, from which we can determine the buyer's position in each marketplace. In so doing, we can determine the buyer's joint positions. Specifically, we use a $|\Delta|$ -tuple $\bar{x} = \langle x_1, \dots, x_{|\Delta|} \rangle \in \mathcal{X}$ to represent the number of buyers choosing specific actions, where x_i is the number of buyers choosing action δ_i^b , \mathcal{X} is the set of all such possible tuples and we have $\sum_{i=1}^{|\Delta|} x_i = B - 1$ (note that we need to exclude the buyer for which we are calculating the expected utility). The probability of exactly x_i buyers choosing action δ_i^b is $(\omega_i^b)^{x_i}$, and then the probability of this tuple appearing is:

$$\rho^b(\bar{x}) = \binom{B-1}{x_1, \dots, x_{|\Delta|}} \times \prod_{i=1}^{|\Delta|} (\omega_i^b)^{x_i} \quad (2)$$

Now for \bar{x} , we determine the buyer's position in each marketplace as follows. Firstly, we compute the number of other buyers whose bids are greater than the buyer's bid in marketplace m , d_m^b , which is given by:

$$X_m^>(\bar{x}, d_m^b) = \sum_{\delta_i^b \in \Delta: d_{im}^b > d_m^b} x_i \quad (3)$$

where d_{im}^b is the bid placed in marketplace m through action δ_i^b . Similarly, we use $X_m^=(\bar{x}, d_m^b)$ to represent the number of buyers whose bids are equal to the buyer's bid in marketplace m (excluding the buyer itself):

$$X_m^=(\bar{x}, d_m^b) = \sum_{\delta_i^b \in \Delta: d_{im}^b = d_m^b} x_i \quad (4)$$

As there are discrete bids and given $X_m^>(\bar{x}, d_m^b)$ buyers bidding higher than the buyer's bid d_m^b and $X_m^=(\bar{x}, d_m^b)$ buyers bidding equal to d_m^b , the buyer's position in marketplace m could be anywhere from $X_m^>(\bar{x}, d_m^b) + 1$ to $X_m^>(\bar{x}, d_m^b) + X_m^=(\bar{x}, d_m^b) + 1$. This constitutes the buyer's position range in this marketplace. Since $X_m^=(\bar{x}, d_m^b) + 1$ buyers have the same bid, as noted previously, a tie-breaking rule is needed to determine the buyer's position. Here, we adopt a standard rule where each of these possible positions³ occurs with equal probability, i.e. $1/(X_m^=(\bar{x}, d_m^b) + 1)$. Now, given the buyer's position ranges in different marketplaces, we can obtain the set of all possible joint positions for the buyer. Specifically, we use an M -tuple $\bar{v}_x = \langle v_1, \dots, v_M \rangle \in \mathcal{V}_{\bar{x}}$ to represent one of the possible joint positions where

³They are $X_m^>(\bar{x}, d_m^b) + 1, X_m^>(\bar{x}, d_m^b) + 2, \dots, X_m^>(\bar{x}, d_m^b) + X_m^=(\bar{x}, d_m^b) + 1$.

v_m is the buyer's position in marketplace m , and $\mathcal{V}_{\bar{x}}$ is the set of all possible joint positions satisfying the condition $X_m^>(\bar{x}, d_m^b) + 1 \leq v_m \leq X_m^>(\bar{x}, d_m^b) + X_m^=(\bar{x}, d_m^b) + 1$ ($m = 1, \dots, M$). The probability of the buyer having the joint positions $\bar{v}_{\bar{x}}$ given the tuple \bar{x} is:

$$\Phi(\bar{v}_{\bar{x}}) = \prod_{m=1}^M \frac{1}{X_m^=(\bar{x}, d_m^b) + 1} \quad (5)$$

Note that tie-breaking occurs independently for each marketplace. Now we have computed the buyer's joint positions across multiple marketplaces when taking action δ^b , and also calculated the probability of the joint positions occurring. The joint positions will be used when matching the buyer with sellers in marketplaces.

The Action Choices of Sellers. In addition to depending on the positions in different marketplaces, the buyer's expected valuation also depends on the sellers' action choices. Specifically, we use a $|\Delta|$ -tuple $\bar{y} = \langle y_1, \dots, y_{|\Delta|} \rangle \in \mathcal{Y}$ to represent the number of sellers choosing different actions, where y_i is the number of sellers choosing action δ_i^s , and \mathcal{Y} is the set of all such possible tuples and we have $\sum_{i=1}^{|\Delta|} y_i = S$. The probability of this tuple appearing is:

$$\rho^s(\bar{y}) = \binom{S}{y_1, \dots, y_{|\Delta|}} \times \prod_{i=1}^{|\Delta|} (\omega_i^s)^{y_i} \quad (6)$$

The Expected Valuation. Given the buyer's joint positions $\bar{v}_{\bar{x}}$ and the number of sellers choosing different actions \bar{y} , we are ready to calculate the buyer's expected valuation for traded goods. Since the buyer can enter multiple marketplaces and thus purchase multiple goods, we need to compute its expected valuation for different units of goods. Recall we assume each trader can only trade one unit of good in each marketplace, and thus when there are M marketplaces in total, the possible number of goods the buyer can purchase is from 1 to M . Specifically, in Section 3.1, we defined the buyer's valuation $v^b(\theta^b, T)$ on T units of goods by considering different types of goods (see Equation 1). Now by considering all possible subsets of marketplaces with cardinality T , where this buyer makes a transaction in each of them, we compute the buyer's expected valuation when it purchases T units of goods given its joint positions $\bar{v}_{\bar{x}}$ and the number of sellers choosing different actions \bar{y} :

$$\begin{aligned} \tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) &= \sum_{\mathcal{M}_T \subset 2^{\mathcal{M}}: |\mathcal{M}_T|=T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_T) \times v^b(\theta^b, T) \\ &= \sum_{\mathcal{M}_T \subset 2^{\mathcal{M}}: |\mathcal{M}_T|=T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_T) \times \alpha_T^b \times \theta^b \end{aligned} \quad (7)$$

where $\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I)$ indicates whether the buyer makes transactions in the subset of marketplaces \mathcal{M}_I and does not make transactions in the remaining marketplaces $\mathcal{M} - \mathcal{M}_I$. Note that given the buyer's joint positions $\bar{v}_{\bar{x}}$ and the number of sellers choosing different actions \bar{y} , whether the buyer is transacting or not in each individual marketplace is independent. Thus whether the buyer transacts in \mathcal{M}_I and not in $\mathcal{M} - \mathcal{M}_I$ is given by:

$$\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) = \prod_{m \in \mathcal{M}_I} \psi^b(v_m, \bar{y}, m, d_m^b) \times \prod_{m \in \mathcal{M} - \mathcal{M}_I} \chi^b(v_m, \bar{y}, m, d_m^b) \quad (8)$$

where $\psi^b(v_m, \bar{y}, m, d_m^b)$ indicates whether the buyer with bid d_m^b makes a transaction in marketplace m given its position v_m and the number of sellers choosing different actions \bar{y} , and $\chi^b(v_m, \bar{y}, m, d_m^b)$ indicates whether the buyer with bid d_m^b does not transact in marketplace m .

We derive $\psi^b(v_m, \bar{y}, m, d_m^b)$ and $\chi^b(v_m, \bar{y}, m, d_m^b)$ as follows. Since the marketplace uses equilibrium matching to match buyers with sellers, if the buyer can transact in marketplace m , then the number of asks which are not greater than its bid, d_m^b , will not be less than its position v_m . Given the number of sellers choosing different actions, we know what asks are placed in marketplace m , from which we can calculate the number of asks not greater than d_m^b :

$$Y_m^{\leq}(\bar{y}, d_m^b) = \sum_{\delta_i^s \in \Delta: d_{im}^s \leq d_m^b} y_i \quad (9)$$

Now, whether the buyer transacts in marketplace m is given by:

$$\psi^b(v_m, \bar{y}, m, d_m^b) = \begin{cases} 1 & \text{if } Y_m^{\leq}(\bar{y}, d_m^b) \geq v_m \\ 0 & \text{if } Y_m^{\leq}(\bar{y}, d_m^b) < v_m \end{cases} \quad (10)$$

and whether the buyer does not transact in m is:

$$\chi^b(v_m, \bar{y}, m, d_m^b) = 1 - \psi^b(v_m, \bar{y}, m, d_m^b) \quad (11)$$

Finally, by considering all possible numbers of units the buyer is purchasing, all possible numbers of sellers choosing different actions, all possible joint positions and all possible numbers of buyers choosing different actions, the buyer's expected valuation is given by:

$$\tilde{V}(\theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{\bar{x} \in \mathcal{X}} \rho^b(\bar{x}) \times \sum_{\bar{v}_{\bar{x}} \in \mathcal{V}_{\bar{x}}} \Phi(\bar{v}_{\bar{x}}) \times \sum_{\bar{y} \in \mathcal{Y}} \rho^s(\bar{y}) \times \sum_{T=1}^M \tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) \quad (12)$$

3.2.2. The Expected Payment

After deriving the expected valuation, in the following, we derive the expected payment that the buyer will make for the goods given the action distributions of buyers and sellers, Ω^b and Ω^s . Firstly, we derive the buyer's expected payment given its joint positions $\bar{v}_{\bar{x}}$ and the specific numbers of sellers choosing different actions \bar{y} . This is equal to the sum of the expected payments in each marketplace, which is:

$$\tilde{\mathcal{P}}^b(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{m=1}^M \tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s) \quad (13)$$

where $\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s)$ is the expected payment of the buyer given its bid d_m^b and its position v_m in marketplace m . Specifically, when $d_m^b = \ominus$, i.e. not bidding in this marketplace, the expected payment is 0; when $d_m^b \neq \ominus$, by sorting the asks ascendingly, we will know what exact ask will be (if it can be) matched with bid d_m^b . We denote the ask matched with bid d_m^b in marketplace m as d_m^s . As a result, $\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s)$ is given by:

$$\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s) = \begin{cases} 0 & \text{if } d_m^b = \ominus \\ d_m^s \times k_m + d_m^b \times (1 - k_m) + \iota & \text{if } d_m^b \geq d_m^s \\ \iota & \text{otherwise} \end{cases} \quad (14)$$

where ι is the constant cost.

Now by considering all the possible numbers of sellers choosing different actions, all possible joint positions and all possible numbers of buyers choosing different actions, the buyer's expected payment is given by:

$$\tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{\bar{x} \in \mathcal{X}} \rho^b(\bar{x}) \times \sum_{\bar{v}_{\bar{x}} \in \mathcal{V}_{\bar{x}}} \Phi(\bar{v}_{\bar{x}}) \times \sum_{\bar{y} \in \mathcal{Y}} \rho^s(\bar{y}) \times \tilde{\mathcal{P}}^b(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s) \quad (15)$$

Finally, the expected utility of the buyer with type θ^b using action δ^b is:

$$\tilde{U}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) = \tilde{V}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) - \tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) \quad (16)$$

3.3. Equilibrium Strategies

After deriving the traders' expected utilities when trading across multiple double auction marketplaces, we are now ready to define the trader's Bayes-Nash equilibrium trading strategy. Since we cannot run the fictitious play algorithm for an infinite number of iterations, we adopt the standard approach of running the algorithm until we obtain an ϵ -Bayes-Nash equilibrium, where each trader cannot

gain more than ϵ utility by unilaterally deviating from its strategy. Formally, the ϵ -Bayes-Nash equilibrium trading strategy is defined as follows:

Definition *The trading strategies $\sigma^b : [0, 1] \rightarrow \Delta$ and $\sigma^s : [0, 1] \rightarrow \Delta$ constitute an ϵ -Bayes-Nash equilibrium if :*

$$\forall i \in \mathcal{B}, \forall \theta^b \in [0, 1], \forall \delta^{b'} \in \Delta$$

$$\tilde{U}_i^b(\theta^b, \sigma^b(\theta^b), \Omega^b, \Omega^s) + \epsilon \geq \tilde{U}_i^b(\theta^b, \delta^{b'}, \Omega^b, \Omega^s)$$

$$\text{and } \forall j \in \mathcal{S}, \forall \theta^s \in [0, 1], \forall \delta^{s'} \in \Delta$$

$$\tilde{U}_j^s(\theta^s, \sigma^s(\theta^s), \Omega^b, \Omega^s) + \epsilon \geq \tilde{U}_j^s(\theta^s, \delta^{s'}, \Omega^b, \Omega^s)$$

where Ω^b and Ω^s are the action distributions of buyers and sellers when they adopt the trading strategy σ^b and σ^s .

4. The Fictitious Play Algorithm

In this section, we describe how to use fictitious play (FP) to approximate a pure Bayes-Nash equilibrium market selection and bidding strategy for traders using the equations derived in Section 3.2. Note that, in our setting, the trader's action actually combines the selected marketplaces and the offers in the selected marketplaces as a whole. Then, in the fictitious play iteration for choosing the best action, traders actually make decisions about which marketplace or marketplaces to enter for all marketplaces simultaneously. While the learning process is dynamic, our aim is to find the equilibrium of this one-shot game.

In what follows, we first briefly describe the standard FP algorithm. We then describe how to compute the trader's best response action (which needs to maximise its expected utility as derived in Section 3.2) against current FP beliefs, and describe how to update FP beliefs according to the best response action distributions. Furthermore, we introduce how to measure the convergence in our setting. Finally, we show the structure of the entire algorithm.

4.1. Standard Fictitious Play

In the standard FP algorithm (von Neumann and Brown, 1950; Brown, 1951), opponents are assumed to play a mixed strategy. Then, by observing relative frequencies of different actions, the player can estimate their opponents' mixed strategies, and take a best response to those strategies. 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 a pure strategy, which means that after a number of iterations, the best response strategy of each player is stable. Then all players' best response strategies constitute a pure Nash equilibrium. Second, it may converge in FP beliefs. Then the converged FP beliefs constitute a mixed Nash equilibrium. We should note that, in reality, it is impossible to run the algorithm to convergence since it involves an infinite number of iteration rounds. Therefore, it is often used to approximate the Nash equilibrium (i.e. deriving the ϵ -Nash equilibrium) by running the fictitious play algorithm for a finite number of rounds.

However, the standard FP algorithm is not suitable for analysing Bayesian games in which there is incomplete information about traders' types. In such games, a strategy is a function that maps the set of player types to the set of allowed actions for the player. 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 address this, Rabinovich et al. (2013) provided a generalised fictitious play algorithm to analyse Bayesian games with continuous types and a finite action space. Using this algorithm, when the FP beliefs converge, they either directly converge to a pure Bayes-Nash equilibrium, or can be purified to produce a pure Bayes-Nash equilibrium (Radner and Rosenthal, 1982). Moreover, it is known that a pure Bayes-Nash equilibrium always exists given the conditions that the game is non-atomic (giving zero probability to any specific player type appearing), and that the action space is finite. However, Rabinovich et al. (2013) only showed how to use this algorithm to analyse traders' strategies in single-sided auctions. Building on this, we extend their fictitious play algorithm to approximate the Bayes-Nash equilibrium market selection and bidding strategies for traders in the much more complex environment of multiple double auction marketplaces.

4.2. Computing the Best Response

We now describe how to compute the best response actions against current FP beliefs. Recall that in Section 3.1, 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 compute the buyers' best response function as follows (mutatis mutandis for sellers):

$$\sigma^{b*}(\theta^b|\Omega^b, \Omega^s) = \operatorname{argmax}_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s) \quad (17)$$




Figure 1: Piecewise linear expected utility functions.

where σ^{b*} is the best response action of the buyer with type θ^b given FP beliefs Ω^b and Ω^s . The optimal utility that a buyer with type θ^b can achieve is:

$$\tilde{U}^*(\theta^b, \Omega^b, \Omega^s) = \max_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s) \quad (18)$$

From the buyer's expected utility equations (see Equations 7 and 12) 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. Given this, and given a finite number of actions, the best response 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 Θ (meaning no participation in the marketplace), 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 $_{\Theta}$* (i.e. x-axis) respectively. The optimal utility achieved by the buyer is represented by the set of thick piecewise linear segments. Note that each line segment corresponds to a set of types on the x-axis, and this actually indicates the buyers' optimal trading strategy given current FP beliefs. Here, the set of types that have a best response action δ_i^b is given by $\sigma^{-b}(\delta_i^b)$ ($i = 1, 2, 3$) and the type set corresponding to the best response action Θ is $\sigma^{-b}(\Theta)$.

4.3. Updating Fictitious Play Beliefs

We have computed the best response function and also provided the sets of types corresponding to the best response actions. Based on this, we can now calculate the best response action distribution of buyers. Given that the upper envelope is a piece-wise linear function, we know that the set of buyer types corresponding to the best response action δ_i^b is $\sigma^{-b}(\delta_i^b)$. Given this, the probability of action δ_i^b

being played by a buyer is given by:

$$\omega_i^b = \int_{x \in \sigma^{-b}(\delta_i^b)} f^b(x) dx \quad (19)$$

By calculating the probability of each action being used, we obtain the current best response action distribution of buyers, denoted by Ω_{br}^b , which is given current FP beliefs. We can then update the FP beliefs of buyers' actions:

$$\Omega_{\tau+1}^b = \frac{\tau}{\tau+1} \times \Omega_{\tau}^b + \frac{1}{\tau+1} \times \Omega_{br}^b \quad (20)$$

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 τ , and Ω_{br}^b is the probability distribution of the best response actions against FP beliefs Ω_{τ}^b . This equation actually gives the FP beliefs at the current round as the average of the best response action distributions of buyers in all previous rounds. 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.

4.4. Measuring Convergence

We now describe how to check the convergence to a Bayes-Nash equilibrium. In the above, we have given the definition of the ϵ -Bayes-Nash equilibrium where each trader cannot gain more than ϵ in expected utility by unilaterally deviating from its strategy. According to this definition, if the trader cannot gain more than ϵ by taking a best response action against the current best response action distributions, the FP algorithm stops the iteration process. At this point, the current best response actions with corresponding type sets constitute an ϵ -Bayes-Nash equilibrium. Formally, the measure of convergence is given by:

$$|\tilde{U}^b(\Omega_{br}^b, \Omega_{br}^s) - \tilde{U}_{br}^b(\Omega_{br}^b, \Omega_{br}^s)| \leq \epsilon \text{ and } |\tilde{U}^s(\Omega_{br}^b, \Omega_{br}^s) - \tilde{U}_{br}^s(\Omega_{br}^b, \Omega_{br}^s)| \leq \epsilon \quad (21)$$

where $\tilde{U}^b(\Omega_{br}^b, \Omega_{br}^s)$ is the expected utility of a buyer in the best response action distributions Ω_{br}^b and Ω_{br}^s :

$$\tilde{U}^b(\Omega_{br}^b, \Omega_{br}^s) = \int_0^1 f^b(x) \times \tilde{U}^b(x, \delta^b, \Omega_{br}^b, \Omega_{br}^s) dx \quad (22)$$

where δ^b is the action chosen by the buyer with type x (actually, it is the best response action of this buyer against FP beliefs Ω_{τ}^b and Ω_{τ}^s , i.e. $\sigma^{b*}(x|\Omega_{\tau}^b, \Omega_{\tau}^s)$).

$\tilde{U}_{br}^b(\Omega_{br}^b, \Omega_{br}^s)$ is the expected utility of a buyer adopting the best response action against the current best response action distributions Ω_{br}^b and Ω_{br}^s :

$$\tilde{U}_{br}^b(\Omega_{br}^b, \Omega_{br}^s) = \int_0^1 f^b(x) \times \tilde{U}^b(x, \delta^{b*}, \Omega_{br}^b, \Omega_{br}^s) dx \quad (23)$$

where $\delta^{b*} = \sigma^{b*}(x|\Omega_{br}^b, \Omega_{br}^s)$ is the best response action of the buyer with type x given action distributions Ω_{br}^b and Ω_{br}^s . The equations for sellers are analogous.

4.5. Algorithm Overview

Given the calculation of best response actions, the update of FP beliefs and the measure of convergence, Figure 2 shows the structure of the entire FP algorithm.

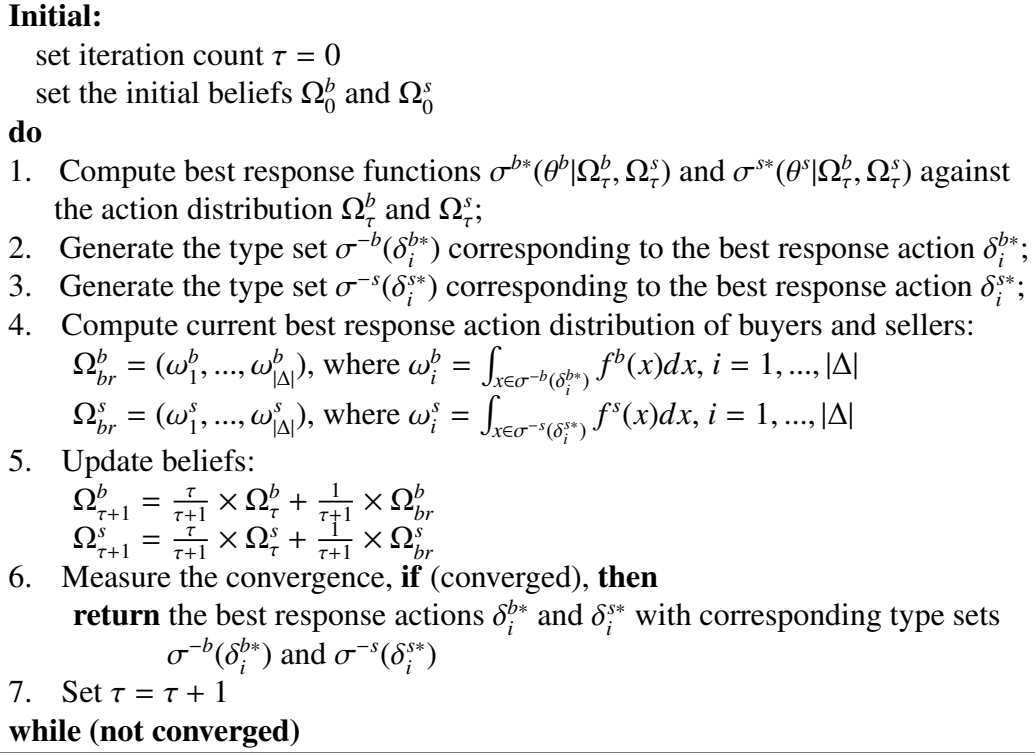


Figure 2: The fictitious play algorithm.

5. Equilibrium Analysis of Market Selections and Bidding Strategies

In this section, we use the FP algorithm from Section 4 to analyse the traders' equilibrium strategies. We first introduce the experimental setting in this paper. After this, we report the convergence results. We then analyse the traders' equilibrium bidding strategies in an isolated marketplace. This helps us to understand the analysis in the more complex setting with multiple marketplaces.

5.1. Experimental Settings

In each of the following experiments, we consider 200 instantiations with different initial FP beliefs on the traders' actions. Each initial belief is uniformly chosen between 0 and 1, and then scaled so they sum to 1. The remaining experimental settings are given in Table 1.⁴ Specifically, we consider 2 marketplaces, 5 buyers and 5 sellers, and 11 discrete bid(ask) levels plus \emptyset (denoting the action when the marketplace is not chosen). Furthermore, we assume that both buyer and seller types are independently drawn from a uniform distribution.⁵ In addition, we set $\epsilon = 0.00001$ in the ϵ -Bayes-Nash equilibrium, and assume that the small cost for traders to enter a marketplace is set to $\iota = 0.0001$. For the transaction price, we assume that $k_m = 0.5$, i.e. the transaction price is set in the middle of the matched bid and ask, which means that the marketplace has no bias in favour of buyers or sellers when allocating profits.

5.2. Convergence Results

As per Section 5.1, we consider 200 different initial FP beliefs on the traders' actions for each experiment. Our experiments empirically demonstrate that, in this setting, the FP algorithm consistently converges.⁶ This allows the identification of the pure Bayes-Nash equilibrium market selection and bidding strategies. Furthermore, we find that the FP algorithm converges to the same Nash equilibrium

⁴Note that the exact Nash equilibrium solution depends on the specific experimental setting. However, although different experimental parameter values result in different equilibrium solutions, the insights drawn from the solutions are broadly similar. For example, as we show in Appendix A, the insights obtained in the setting with three marketplaces are similar to those with two.

⁵In our setting, it is easy to consider different distributions. We just need to set the probability density function $f^b(x)$ in Equations 19, 22 and 23 to any specific probability density function for buyers, and do this for sellers analogously.

⁶Note that, while a pure-strategy Nash equilibrium always exists in this setting (Rabinovich et al., 2013), in general, the FP algorithm is not guaranteed to converge (Vijay and Sjöström, 1998).

Parameter	Value(s)
Number of marketplaces M	2
Number of buyers B	5
Number of sellers S	5
Trader offer space Ψ	$\{0, 0.1, \dots, 0.9, 1, \ominus\}$
Buyer type probability distribution F^b	Uniform($[0, 1]$)
Seller type probability distribution F^s	Uniform($[0, 1]$)
Buyer preference coefficient function $\alpha^b(T)$	varies
ϵ value in ϵ -Bayes-Nash equilibrium	$\epsilon = 0.00001$
Constant cost ι	$\iota = 0.0001$
k value in k -pricing policy	$k = 0.5$

Table 1: Experimental setting

(see Figures 3 to 11). The only exception is a setting where the initial FP beliefs are such that no buyers and sellers participate in the marketplace. In this case, the algorithm will immediately converge to a different Bayes-Nash equilibrium where nobody enters the marketplace. However, this equilibrium is trivial since there is a zero probability of transactions happening between buyers and sellers. This trivial equilibrium exists both in isolated marketplaces and when trading across multiple double auctions where no buyers and sellers enter any marketplaces. Given this, in what follows we focus solely on non-trivial equilibria where there is a positive probability of transactions happening.

Moreover, for each of the following experiments, the average number of iterations to convergence and standard deviation are shown in Table 2. Note that, although each experiment is run 200 times with different initial beliefs, the results shown in Table 2 are the average over the 150 runs with the lowest iteration numbers to convergence. We do this because the iteration numbers to convergence vary depending on the initial beliefs. Although in most cases the algorithm converges quickly, there are a few initial FP beliefs which take much longer to converge (sometimes by a factor of 6) and including these would skew the results. We avoid this problem by removing the top 50 runs. From the experiments, we can see that, as the number of traders increases (e.g. comparing experiment 1 with 5 buyers and 5 sellers to experiment 2 with 8 buyers and 5 sellers), the average number of iterations to convergence increases. We furthermore find that, when traders bid asymmetrically in two marketplaces (e.g. experiments 5, 8 and 9 shown in Figures 5, 8 and 9 respectively), the average number of iterations to convergence also increases.

Experiment	Average Number	Standard Deviation
1. An isolated marketplace	604.26	235.40
2. An isolated marketplace, 8 buyers and 5 sellers	1084.38	384.31
3. Single-home	654.46	241.53
4. Multi-home with independent goods	752.94	282.58
5. Multi-home with perfectly substitutable goods	1348.32	493.24
6. Multi-home with perfectly complementary goods	682.56	247.23
7. Multi-home with complementary goods, $\alpha^b(1) = 1$ and $\alpha^b(2) = 4$	694.40	254.45
8. Hybrid with independent goods	1638.98	608.60
9. Hybrid with substitutable goods, $\alpha^b(1) = 1$ and $\alpha^b(2) = 1.8$	1758.36	653.11
10. Hybrid with perfectly complementary goods, 2 buyers and 10 sellers	913.52	341.02
11. Hybrid with complementary goods, $\alpha^b(1) = 1$ and $\alpha^b(2) = 4$	791.26	284.82

Table 2: Average number of iterations to convergence and standard deviation for each experiment where parameter values are taken from Table 1, unless mentioned otherwise.

5.3. Trading in an Isolated Marketplace

We now analyse the traders' Nash equilibrium strategies. As introduced in Section 2, there are many heuristic bidding strategies in double auctions. However, they all fail to determine exactly what traders should bid in equilibrium. This is an important shortcoming since the way in which traders bid in a given marketplace affects their expected utilities and this, in turn, affects their selection of marketplaces.

The converged pure-strategy Bayes-Nash equilibrium bidding strategies from the FP algorithm are shown in Figure 3. The gray line represents the buyers' bids in equilibrium for given types and the black line represents the sellers' asks in equilibrium. We find that buyers shade their bids (i.e. bid less than their types) and sellers shade their asks (i.e. ask more than their types). They do so in order to keep profits. 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 small cost ι .

In the above, we considered the case with the same number of buyers and sellers. We now analyse asymmetric settings where these numbers differ. In particular, Figure 4 shows the results for the case with 8 buyers and 5 sellers. As 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. In this environment, sellers have more market power, and can therefore extract more profit from the transactions.




Figure 3: Equilibrium strategies of traders in an isolated marketplace with 5 buyers and 5 sellers.

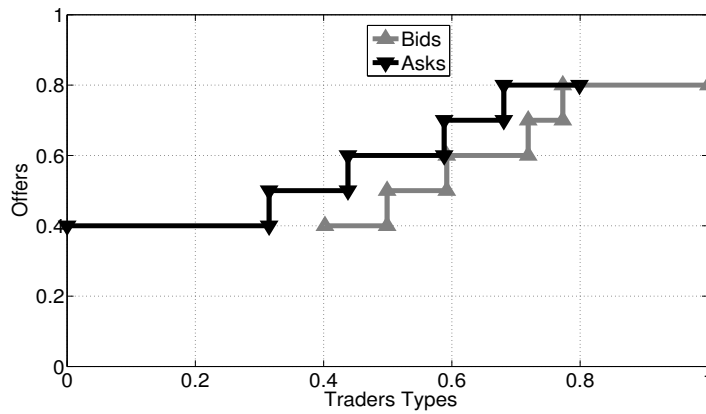


Figure 4: Equilibrium strategies of traders in an isolated marketplace with 8 buyers and 5 sellers.

5.4. Trading Across Multiple Marketplaces

The above analysis of isolated marketplaces is important for understanding the setting where traders trade across multiple marketplaces. For example, as we will show, in the setting with multiple marketplaces, when traders converge to one marketplace, their bidding strategies are exactly the same as those in an isolated marketplace. In the following, we start by analysing equilibrium market selection and bidding strategies for traders in environments with multiple marketplaces. As mentioned in Section 1, in addition to the isolated marketplace setting, there are three other types of trading environments: single-home, multi-home and hybrid. Furthermore, when multiple goods are traded, goods can be either independent, substitutable or complementary. In the single-home trading environment, since

we assume that only one unit of the good is traded by each trader, we do not need to consider the types of goods. However, in the multi-home and hybrid environments where multiple goods can be traded by each trader, we do need to consider different types of goods. In the following, we first analyse the traders' equilibrium strategies in the single-home environment. Then we extend the analysis to the multi-home and hybrid environments by considering different types of goods.

5.4.1. *Single-Home Trading*


We now consider traders' equilibrium strategies in the single-home environment with two marketplaces, where traders can only enter one marketplace at a time. In this environment, traders can only trade one unit of the good. Thus the valuation for the good is equal to its type, i.e. $\alpha^b(1) = \alpha^s(1) = 1$. By using FP, we find that, except for some traders (buyers with low types and sellers with high types) choosing no marketplace, all other traders eventually converge to one marketplace in equilibrium. Since the two marketplaces are identical in this environment, 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 those in an isolated marketplace (i.e. Figure 3).

When considering the traders' equilibrium strategies across more than two marketplaces, we still obtain similar conclusions. Specifically, traders converge to one marketplace, and their bidding strategies in this marketplace are the same as in Figure 3.

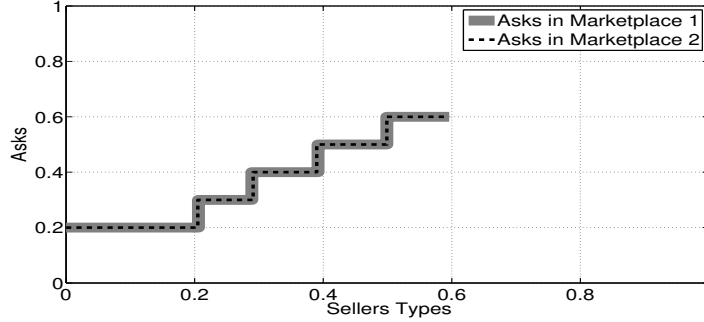
5.4.2. *Multi-Home Trading*

So far we have analysed the equilibrium strategies of traders in single-home trading environments. Now we extend the analysis to the multi-home settings where both buyers and sellers can enter multiple marketplaces at a time. In such environments, multiple goods may be traded by each trader. Therefore, we now need to consider the types of goods since these will affect traders' strategies. Specifically, here we assume that, for sellers, all goods are independent since a seller's valuation for sold goods is usually equal to the sum of its valuation for each individual sold good (i.e. additive). Thus, in the multi-home environment, sellers are willing to enter any marketplace which can provide positive profits for them. For buyers, goods can be either independent, substitutable or complementary. In the following, we will analyse the equilibrium strategies for traders by considering different types of goods for buyers.

Independent Goods. In the multi-home environment with independent goods, buyers' values on goods are additive, i.e. the buyer preference coefficient $\alpha^b(1) =$



(a) Buyers



(b) Sellers

Figure 5: Equilibrium strategies of traders in the multi-home trading environment with perfectly substitutable goods for buyers, 5 buyers, 5 sellers and 2 marketplaces.

1 and $\alpha^b(2) = 2$. In such situations, buyers and sellers will enter both marketplaces when their profits in both marketplaces are positive, and then their equilibrium bidding strategies are exactly the same as those in the isolated marketplace setting (see Figure 3).

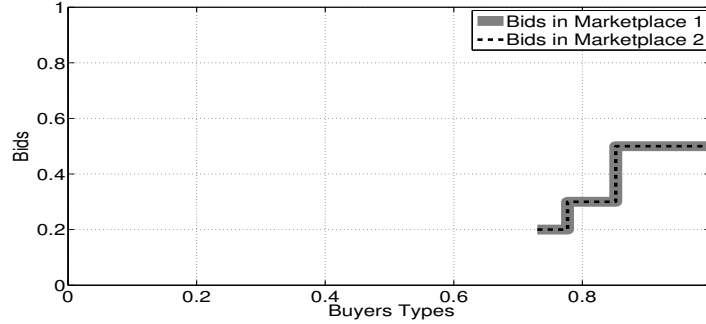
Substitutable Goods. We now consider the equilibrium strategies for traders with substitutable goods for buyers (i.e. $\alpha^b(1) = 1$, and $1 \leq \alpha^b(2) < 2$).

Firstly, we consider perfectly substitutable goods, i.e. $\alpha^b(1) = 1$ and $\alpha^b(2) = 1$. This means that, when the buyer with type θ^b wins one good, it obtains valuation θ^b and pays for the good; and when this buyer wins two goods, it only obtains valuation θ^b , but pays for two goods. The results are shown in Figures 5(a) and 5(b). We find that, in this setting, even though the marketplaces are identical, a buyer with a particular type will bid differently in the two marketplaces. The implications are as follows. Because of perfectly substitutable goods, buyers would


like to purchase and pay for only one good. Therefore, buyers with high types choose to only bid in one marketplace, and bid slightly higher than the buyers with lower types. By so doing, they can win one good successfully and only pay for this good. For example, when the types are within $[0.928, 1.0]$, they bid 0.6 in marketplace 1. However, buyers with lower types have to bid in both marketplaces in order to increase the probability of being matched. In this situation, they may transact in both marketplaces, which means they have to pay for two goods. Then, in order to maximise the profit of the transaction and minimise the payment, they may choose to bid slightly higher in one marketplace and lower in another. Furthermore, we note that some of these buyers with lower types will bid higher in one marketplace, and others will bid higher in the other one. This is because all buyers bidding higher in one marketplace results in fierce competition between them. Thus some of them will choose to bid higher in another marketplace in order to increase the probability of being matched. This result provides useful insights for buyers trading perfectly substitutable goods. Specifically, when buyers' valuations on individual goods are high, they only need to bid in one marketplace. However, if their valuations are low, they need to bid in multiple marketplaces to increase their probability of transacting.

Furthermore, we extend this analysis by increasing the buyer's valuation for the substitutable goods (i.e. $\alpha^b(2)$ is increased from 1 to 2). In so doing, the goods become less substitutable (no longer perfectly substitutable). We find that, as the goods become less substitutable, buyers with high types begin to bid in two marketplaces since they start to prefer to win multiple goods, and bidding in both marketplaces increases the probability of being matched. Note that when the (perfectly) substitutable goods are traded, the exposure problem may arise. This occurs when a buyer intends to win a certain number of items by placing bids in all marketplaces to increase the probability of transacting. In this case a buyer could end up with with a loss when its overall payment exceeds its overall valuation for the obtained items.

Complementary Goods. Now we analyse the setting with complementary goods for buyers. We first consider the case with perfectly complementary goods, i.e. $\alpha^b(1) = 0$ and $\alpha^b(2) = 1$. This means that when the buyer with type θ^b wins one good, it obtains zero valuation, but needs to pay for the good; and when this buyer wins two goods, it obtains valuation θ^b , and pays for two goods. The results are shown in Figures 6(a) and 6(b). We can see that, with perfectly complementary goods, only buyers with high types bid in marketplaces, and they bid in both marketplaces (if they only bid in one marketplace and transact, they need to pay for the



(a) Buyers




(b) Sellers

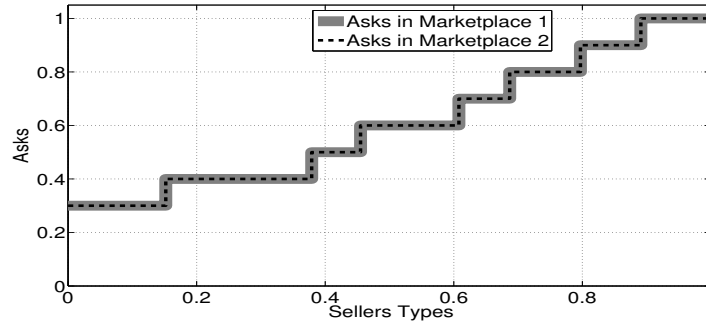
Figure 6: Equilibrium strategies of traders in the multi-home trading environment with perfectly complementary goods for buyers, 5 buyers, 5 sellers and 2 marketplaces.

good, but obtain zero valuation because of the perfectly complementary goods). Buyers with relatively low types do not enter any marketplace since even though they may bid in both marketplaces, because of their low types and perfectly complementary goods, they are more likely to obtain losses. Furthermore, we find that buyers shade their bids more than what they do in the case of isolated marketplaces. For example, when the buyers' types are within $[0.825, 1.0]$, buyers bid 0.5 in both marketplaces. The reason is as follows. With perfectly complementary goods, the buyer with type θ^b purchasing two units, only obtains valuation θ^b , but has to pay for two goods. In order to ensure positive profits, buyers should shade their bids more.

Furthermore, we extend this analysis to the case of non-perfectly complementary goods where $\alpha^b(1) = 1$ and $\alpha^b(2) = 4$. That is, when the buyer with type θ^b wins one good, it obtains valuation θ^b and pays for the good; and when this buyer wins two goods, it obtains valuation $4 \times \theta^b$ (four times its type). The results are



(a) Buyers



(b) Sellers

Figure 7: Equilibrium strategies of traders in the multi-home trading environment with complementary goods for buyers where $\alpha^b(1) = 1$ and $\alpha^b(2) = 4, 5$ buyers, 5 sellers and 2 marketplaces.

shown in Figures 7(a) and 7(b). Compared to the perfectly complementary goods, we find that buyers still bid in two marketplaces. However, they do not shade their bids. Instead, buyers increase their bids in order to increase the probability of being matched. Particularly when buyers' types are within $[0.408, 1.0]$, buyers bid 1.0 in each marketplace. This is because when purchasing two complementary goods, the buyers' values for the goods are very high (i.e. $4 \times \theta^b$). In this situation, in terms of obtaining more profits, it is better for the buyers to increase their bids to increase the probability of transacting than it is to shade. Furthermore, this will cause a larger range of sellers to ask since when buyers increase their bids, sellers with high types can still trade. Note that when the (perfectly) complementary goods are traded, the exposure problem may also arise. This occurs when the buyer, in order to secure multiple complementary goods, bids more than the valuation for a single good. The buyer could then face a loss when failing to win

sufficient goods.

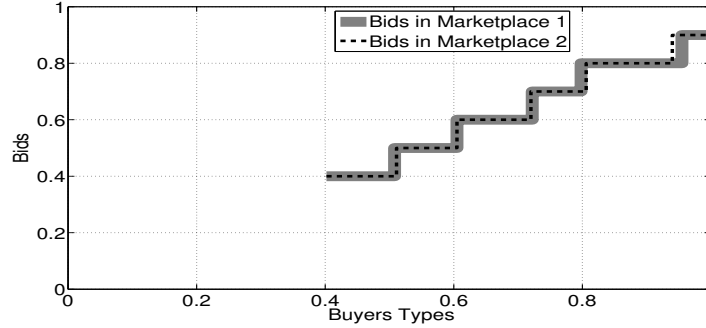
In Conclusion. In the experiments with multi-home environments, we find a number of insights for trading strategies. Specifically, in the case of independent goods, traders can enter both marketplaces when their expected utilities in these marketplaces are positive, and their equilibrium bidding strategies are then identical to those in the isolated marketplace setting. In the multi-home environment with perfectly substitutable goods, buyers with high types should only bid in one marketplace, and buyers with relatively lower types should bid high in one marketplace and bid low in another. Furthermore, as the buyers' values on the substitutable goods increase, buyers with high types begin to bid in two marketplaces. Then in the environment with perfectly complementary goods, buyers need to shade more in order to guarantee positive profits. As buyers' values on multiple complementary goods increase, buyers can increase their bids instead of shading in order to win more goods. Similar insights are found when the analysis is extended to the case with more than two marketplaces (refer to the experiments detailed in Appendix A.2).

5.4.3. Hybrid Trading


We now investigate the equilibrium strategies of traders in the hybrid trading environment where one side can only participate in one marketplace (i.e. single-home trading), and the other side can participate in multiple marketplaces (i.e. multi-home trading). Specifically, we consider the case that buyers can participate in multiple marketplaces and sellers can only participate in one marketplace at a time.⁷ Note that the analysis of the opposite case where sellers can choose multiple marketplaces and buyers can only choose one provides similar results.

Independent Goods. The case with independent goods for buyers (i.e. $\alpha^b(1) = 1$, $\alpha^b(2) = 2$) is shown in Figures 8(a) and 8(b). From these figures, we can see that, in contrast to the analysis in the single-home environment (see Section 5.4.1) where sellers converge to one marketplace in equilibrium, in this case sellers split and place asks in different marketplaces in equilibrium. The reason is as follows. The sellers have to compete with each other in order to be matched with buyers. Thus they prefer marketplaces with fewer sellers. Because both marketplaces have the same buyers in this case, the impact of buyers on the sellers' strategies

⁷Note that in the iteration of fictitious play, the seller can switch between marketplaces to select the best marketplace to maximise its expected utility.



(a) Buyers



(b) Sellers

Figure 8: Equilibrium strategies of traders in the hybrid trading environment with independent goods for buyers, 5 buyers, 5 sellers and 2 marketplaces.

in both marketplaces are identical. In this environment, the internal competition between the sellers becomes the dominant effect and this drives the sellers to stay in different marketplaces. Another interesting and non-obvious result 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 into two marketplaces, in each marketplace the number of sellers is less than the number of buyers. Thus, as per our previous analysis (see Figure 4), sellers have more market power than buyers, and so buyers raise their bids (i.e. shade less) in order to be matched, and sellers raise their asks to extract more profits from the transactions.

Substitutable Goods. Now we turn to the case when goods are substitutable for buyers. We first analyse the case with perfectly substitutable goods (i.e. $\alpha^b(1) = 1$, $\alpha^b(2) = 1$). In this case, we find that in equilibrium traders will only choose one

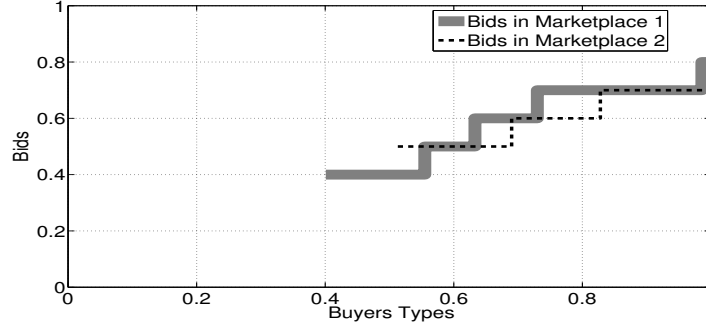
marketplace, and the bidding strategies are the same as the isolated marketplace (see Figure 3). This is because when goods are perfectly substitutable for buyers, buying more than one good does not mean obtaining more values, but means paying more. Thus buyers will prefer to bid in one marketplace. This causes sellers to converge to that marketplace.

Furthermore, since in the above analysis when buyers' expected values of goods are additive (i.e. independent), we find that buyers bid in multiple marketplaces. Thus we hypothesise that as buyers' values on multiple goods increase (i.e. the goods become less substitutable), they begin to prefer to bid in multiple marketplaces, and sellers may split into two marketplaces because of the internal competition between sellers. This is confirmed by running experiments with different values of $\alpha^b(2)$. In particular, we find that when $\alpha^b(2) \geq 1.8$, buyers begin to bid in two marketplaces. Specifically, when $\alpha^b(2) = 1.8$, the traders' equilibrium bidding strategies are shown in Figures 9(a) and 9(b), from which we can see that buyers bid in two marketplaces and sellers split into two marketplaces.

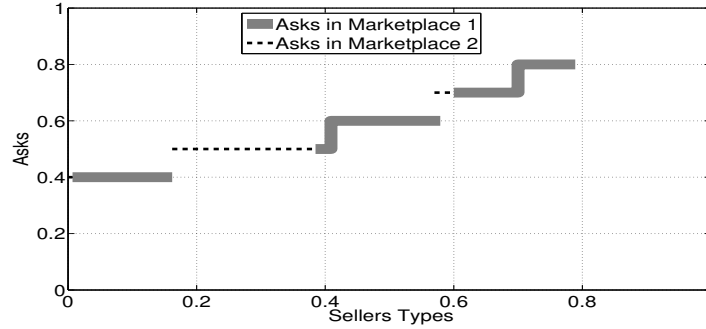
Complementary Goods. Finally, we analyse the case with complementary goods for buyers. Firstly, we consider perfectly complementary goods for buyers (i.e. $\alpha^b(1) = 0$ and $\alpha^b(2) = 1$). When there are 5 buyers and 5 sellers, we find an interesting result where in equilibrium, no traders choose any marketplace. The reason is as follows. With perfectly complementary goods, buyers have to bid in both marketplaces and need to shade their bids in order to make positive profits. However, since sellers adopt single-home trading, sellers have more market power when they split into two marketplaces, and buyers have to increase their bids to increase the probability of being matched. Thus buyers cannot shade too much to win. However, when they purchase two goods, their valuations of the two goods are equal to their types, but they have to pay for both goods. In this situation, buyers will make a loss. Therefore, buyers do not choose any marketplaces, and consequently sellers also choose no marketplaces.

The above situation changes when the market power is changed by amending the number of buyers and sellers, or the buyers' valuation of multiple complementary goods is increased. Specifically, when we increase the number of sellers (to decrease their market power) and decrease the number of buyers (to increase their market power), for example, for 2 buyers and 10 sellers, the results are shown in Figures 10(a) and 10(b). We can see that buyers bid in two marketplaces, and shade more to keep profits. Because of the internal competition, sellers distribute themselves over two marketplaces.

Furthermore, in the case with 5 buyers and 5 sellers, if we change the buyer



(a) Buyers

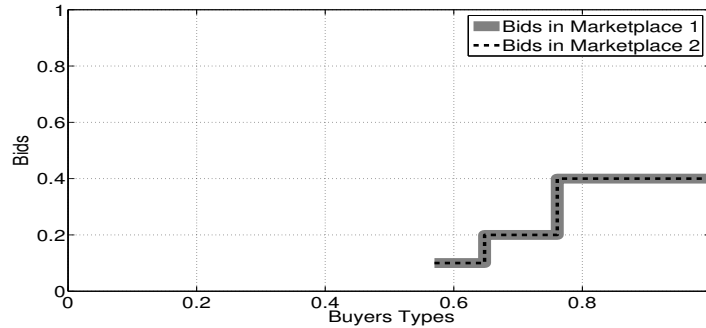


(b) Sellers


Figure 9: Equilibrium strategies of traders in the hybrid trading environment with substitutable goods for buyers where $\alpha^b(1) = 1$ and $\alpha^b(2) = 1.8$, 5 buyers, 5 sellers and 2 marketplaces.

preference coefficient for $\alpha^b(1) = 1$ and $\alpha^b(2) = 4$ (i.e. when purchasing two goods, the buyer with type θ^b can obtain valuation $4 \times \theta^b$), the results are shown in Figures 11(a) and 11(b). We can see that buyers bid in two marketplaces and, instead of shading, they increase their bids to increase their probability of transacting, since they will then obtain very high valuations (i.e. four times their types). For sellers, since the buyers' bids in both marketplaces are identical, each type of seller has an equal probability of asking in each marketplace. For example, when sellers' types are within $[0, 0.257]$, sellers ask 0.6 in either marketplace 1 or marketplace 2. Furthermore, since buyers increase their bids, sellers can increase their asks to obtain more profits.

In Conclusion. In the above experiments, we also find various insights for trading strategies. In particular, in the case with independent goods, sellers should split



(a) Buyers



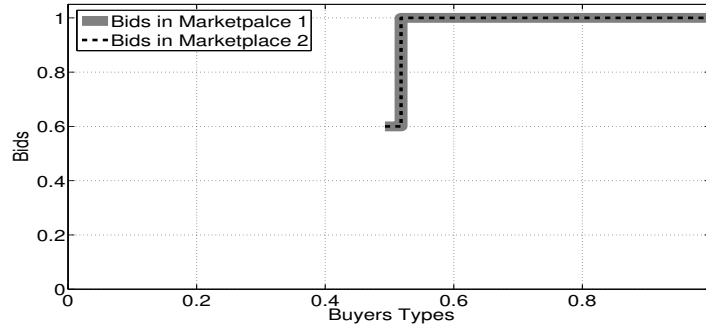
(b) Sellers

Figure 10: Equilibrium strategies of traders in the hybrid trading environment with perfectly complementary goods for buyers, 2 buyers, 10 sellers and 2 marketplaces.


and place asks in both marketplaces. For perfectly substitutable goods, traders converge to bid in one marketplace. As the valuation for the substitutable goods increases, traders begin to bid in both marketplaces. Then for perfectly complementary goods, traders enter no marketplace to avoid the loss. As the buyers' valuation for the complementary goods increases, traders begin to bid in both marketplaces and buyers need to shade less (or even not shade) to increase their probability of transacting. Similar insights can be obtained when we extend the analysis to the case with more than two marketplaces (refer to the experiments in Appendix A.3).

6. Conclusions and Future Work

In this paper, we analysed the Bayes-Nash equilibrium market selection and bidding strategies for traders that operate across multiple double auction market-



(a) Buyers



(b) Sellers

Figure 11: Equilibrium strategies of traders in the hybrid trading environment with complementary goods for buyers where $\alpha^b(1) = 1$ and $\alpha^b(2) = 4, 5$ buyers, 5 sellers and 2 marketplaces.

places in different trading environments with different types of goods. Specifically, we first derived the equations to calculate the expected utilities of traders in this complex setting. Then we used a FP algorithm to analyse the traders' Bayes-Nash equilibrium trading strategies in isolated marketplaces, single-home, multi-home and hybrid trading environments with independent, substitutable and complementary goods respectively.

We first analysed the traders' strategies in isolated marketplaces, where we found that they shade their offers in equilibrium. We then analysed the traders' equilibrium market selection and bidding strategies across multiple double auction marketplaces in single-home trading environments, and found that all traders eventually converge to bid in one marketplace. Moreover, how they bid is exactly the same as that in the isolated marketplace setting.

We extended the analysis by considering multi-home environments with different types of goods. Specifically, we assumed that the trading goods are in-

dependent for sellers, and could be independent, substitutable or complementary for buyers. With independent goods, traders bid in multiple marketplaces and the equilibrium bidding strategies are exactly the same as those in isolated marketplaces. With substitutable goods, we found that buyers with high types only bid in one marketplace and buyers with low types bid in multiple marketplaces to increase their probability of being matched. When perfectly complementary goods are traded, only buyers with high types bid in multiple marketplaces. Buyers with low types do not enter any marketplace. However, when buyers' values on multiple complementary goods increase, more buyers bid in multiple marketplaces and they increase their bids, instead of shading, to win more goods.

Finally, we extended the analysis to hybrid environments where buyers can trade in multiple marketplaces and sellers can only trade in one. We found that, for independent goods, different from the single-home result where sellers converge to bid in one marketplace, sellers split to place asks in different marketplaces in equilibrium. For substitutable goods we found that, as the buyers' values for multiple goods increases, they begin to bid in multiple marketplaces. This causes sellers to distribute themselves over different marketplaces (in contrast to perfectly substitutable goods, all traders only choose one marketplace in equilibrium). Then, for complementary goods, we found that no traders choose any marketplaces when trading perfectly complementary goods and traders begin to participate when the buyers' valuations on complementary goods increase or buyers have more market power than sellers.

In sum, we have provided a number of insights for automated trading agents that operate across multiple double auction marketplaces. However, in order to obtain tractable results, our work makes a number of simplifying assumptions that need to be addressed in the future. The first limitation is that, in order to obtain an equilibrium solution in a reasonable time, we consider a relatively small number of traders and marketplaces. When the number of participants increases beyond 100 agents, the computational load is intense and it is difficult to get the solution in a reasonable time. A possible way to address this limitation is to consider a higher level of abstraction in which individual agents represent many agents of the same type (as per the work of Sandholm (2010)). The second limitation is that our analysis is focused on the clearing house mechanism. In reality, some marketplaces will adopt the continuous double auction mechanism, which means that the marketplace will try and match traders as soon as new offers arrive. In the future, we would like to extend our analysis to cover this mechanism. The third limitation is that we only consider identical goods, and do not consider different goods, especially when they are substitutable or complementary. In the future, we

also intend to address this limitation.

Acknowledgement

This paper was funded by the National Natural Science Foundation of China (No. 61402344), Scientific Research Foundation for the Returned Overseas Chinese Scholars, Ministry of Education of China, and Specialized Research Fund for the Doctoral Program of Higher Education of China (No. 20120092120033).

Appendix A. Supplementary Experiments for Three Marketplaces

We present additional experiments to those in Section 5 by considering three double auction marketplaces. We omit them from the main text because they provide no new and significant insights. They are, nevertheless, useful in providing further support for the claims made in Section 5. Specifically, the parameter setting of the following additional experiments is the same as Table 1, except that we allow three marketplaces.


Appendix A.1. Single-Home Trading

In the single-home trading environment with three double auction marketplaces, we find that traders eventually converge to one marketplace and the equilibrium bidding strategies are the same as those in an isolated marketplace (see Figure 3).


Appendix A.2. Multi-Home Trading

In the multi-home trading environment with independent goods, we find that buyers and sellers enter all three marketplaces when their profits in these marketplaces are positive, and then their equilibrium bidding strategies are exactly the same as those in an isolated marketplace setting (see Figure 3).

The case of perfectly substitutable goods for buyers (i.e. $\alpha^b(1) = 1$, $\alpha^b(2) = 1$ and $\alpha^b(3) = 1$) is shown in Figure A.1. These results are similar to the experiment with 2 marketplaces (see Figure 5). In particular, buyers with high types choose to only bid in one marketplace and bid higher than the buyers with lower types. Compared to Figure 5, we find that buyers with slightly lower types enter two marketplaces, and bid slightly higher in one, and bid lower in the other. For example, when the buyer's type is within $[0.594, 0.632]$, it chooses marketplaces 2 and 3, and submits bids 0.2 and 0.4 respectively. Furthermore, buyers with much lower types choose to bid in three marketplaces to increase the probability of transacting, and bid slightly higher in one marketplace, and bid equally low



(a) Buyers




(b) Sellers

Figure A.1: Equilibrium strategies of traders in the multi-home trading environment with perfectly substitutable goods for buyers, 5 buyers, 5 sellers and 3 marketplaces.


in the other two. For example, when the buyer's type is within $[0.442, 0.496]$, it participates in three marketplaces, and submits bid 0.3 in marketplace 2 and 0.2 in marketplaces 1 and 3.

We now turn to the case of complementary goods for buyers. We first consider the case with perfectly complementary goods, i.e. $\alpha^b(1) = 0$, $\alpha^b(2) = 0$ and $\alpha^b(3) = 1$. This means that when a buyer with type θ^b wins one or two goods, it obtains zero valuation, but needs to pay for the goods; and when it wins three goods, it obtains valuation θ^b , and pays for three goods. The results are shown in Figure A.2. We can see that, with perfectly complementary goods, only buyers with high types bid in marketplaces, and they bid in all three marketplaces. Buyers with relatively low types do not enter any marketplace. Furthermore, compared to Figure 6, because of the difficulty of winning three goods, we find that fewer buyers bid than in the case with two marketplaces.

Furthermore, we extend this analysis to the case of non-perfectly complementary goods where $\alpha^b(1) = 1$, $\alpha^b(2) = 4$ and $\alpha^b(3) = 6$. That is, when the buyer with type θ^b wins one good, it obtains valuation θ^b and pays for the good; and when it wins two goods, it obtains valuation $4 \times \theta^b$ (four times its type); and when



(a) Buyers



(b) Sellers


Figure A.2: Equilibrium strategies of traders in the multi-home trading environment with perfectly complementary goods for buyers, 5 buyers, 5 sellers and 3 marketplaces.

it wins three goods, it obtains valuation $6 \times \theta^b$ (six times its type). The results are shown in Figure A.3 and are similar to the case of two marketplaces (see Figure 7). In particular, buyers do not shade their bids and instead increase their bids in order to increase the probability of being matched.


Appendix A.3. HybridTrading

Finally, we show the supplementary experiments with three marketplaces in the hybrid trading environment. Firstly, we consider the case with independent goods for buyers (i.e. $\alpha^b(1) = 1$, $\alpha^b(2) = 2$ and $\alpha^b(3) = 3$). The results are shown in Figure A.4. Similar to the case with two marketplaces (see Figure 8), we can see that sellers split and place asks in different marketplaces in equilibrium. Furthermore, since sellers distribute themselves over three marketplaces, the number of sellers in each marketplace is less than that with two marketplaces. This means that sellers have more market power. Therefore, compared to Figure 8, sellers shade their offers more and buyers shade their offers less in order to be matched.

We now consider the case of substitutable goods for buyers. With perfectly substitutable goods (i.e. $\alpha^b(1) = 1$, $\alpha^b(2) = 1$ and $\alpha^b(3) = 1$), in equilibrium



(a) Buyers




(b) Sellers


Figure A.3: Equilibrium strategies of traders in the multi-home trading environment with complementary goods for buyers where $\alpha^b(1) = 1$, $\alpha^b(2) = 4$ and $\alpha^b(3) = 6$, 5 buyers, 5 sellers and 3 marketplaces.

traders only choose one marketplace and the bidding strategies are the same as the case with an isolated marketplace (see Figure 3). Furthermore, when buyers' values on multiple goods increase (i.e. the goods become less substitutable), buyers begin to prefer to bid in multiple marketplaces, and sellers may split into three marketplaces because of the internal competition between them. Specifically, when $\alpha^b(1) = 1$, $\alpha^b(2) = 1.8$ and $\alpha^b(3) = 2.5$, the experimental results are shown in Figure A.5. As can be seen, buyers bid in three marketplaces and sellers split into three marketplaces. Furthermore, because there are fewer sellers in each marketplace than the case with two marketplaces, sellers have more market power. Therefore, compared to Figures 9, sellers shade their offers more.

Finally, we consider three marketplaces with complementary goods for buyers. For perfectly complementary goods (i.e. $\alpha^b(1) = 0$, $\alpha^b(2) = 0$ and $\alpha^b(3) = 1$) with 5 buyers and 5 sellers, we observe the same results as two marketplaces. That is, no traders choose any marketplace. However, when there are two buyers and ten sellers (buyers have more market power than sellers), the results are shown in Figure A.6. Here buyers bid in three marketplaces and sellers split into three marketplaces. Furthermore, only when the buyer wins three goods, it obtains valuation θ^b , and needs to pay for three goods. Therefore, compared to the




(a) Buyers




(b) Sellers

Figure A.4: Equilibrium strategies of traders in the hybrid trading environment with independent goods for buyers, 5 buyers, 5 sellers and 3 marketplaces.

case with two marketplaces (see Figure 10), buyers shade more to obtain positive profits. Furthermore, for non-perfectly complementary goods where $\alpha^b(1) = 1$, $\alpha^b(2) = 4$ and $\alpha^b(3) = 6$, the results are shown in Figure A.7. Similar to the case of two marketplaces (see Figure 11), buyers bid in three marketplaces, and instead of shading, they increase their bids to increase the probability of transacting to obtain very high valuations.




(a) Buyers




(b) Sellers

Figure A.5: Equilibrium strategies of traders in the hybrid trading environment with substitutable goods for buyers where $\alpha^b(1) = 1$, $\alpha^b(2) = 1.8$ and $\alpha^b(3) = 2.5$, 5 buyers, 5 sellers and 3 marketplaces.




(a) Buyers




(b) Sellers

Figure A.6: Equilibrium strategies of traders in the hybrid trading environment with perfectly complementary goods for buyers, 2 buyers, 10 sellers and 3 marketplaces.



(a) Buyers



(b) Sellers

Figure A.7: Equilibrium strategies of traders in the hybrid trading environment with complementary goods for buyers where $\alpha^b(1) = 1$, $\alpha^b(2) = 4$ and $\alpha^b(3) = 6$, 5 buyers, 5 sellers and 3 marketplaces.

References

- Anthony, P., Jennings, N. R., 2003. Developing a bidding agent for multiple heterogeneous auctions. *ACM Transactions on Internet Technology* 3 (3), 185–217.
- Anthony, P., Law, E., 2012. Reserve price strategy for seller agent in multiple simultaneous auctions. *International Journal of Knowledge-Based and Intelligent Engineering Systems* 16 (3), 163–175.
- Ashlagi, I., Edelman, B. G., Lee, H. S., 2013. Competing ad auctions. Harvard Business School NOM Unit Working Paper (10-055).
- Bredin, J., Parkes, D. C., 2005. Models for truthful online double auctions. In: *Proceedings of the 21st Conference on Uncertainty in Artificial Intelligence*. pp. 50–59.
- Brown, G. W., 1951. Iterative solution of games by fictitious play. *Activity Analysis of Production and Allocation* 13 (1), 374–376.
- Byde, A., Preist, C., Jennings, N. R., 2002. Decision procedures for multiple auctions. In: *Proceedings of the 1st International Conference on Autonomous Agents and Multi-Agent Systems*. pp. 613–620.
- Cai, K., Niu, J., Parsons, S., 2010. On the economic effects of competition between double auction markets. In: *Agent-Mediated Electronic Commerce and Trading Agent Design and Analysis, Lecture Notes in Business Information Processing*. pp. 88–102.
- Carrie, C., 2008. Illuminating the new dark influence on trading and U.S. market structure. *Journal of Trading* 3(1), 40–55.
- Chatterjee, K., Samuelson, W., 1983. Bargaining under incomplete information. *Operations Research* 31 (5), 835–851.
- Dang, C., Hu, Q., Liu, J., 2015. Bidding strategies in online auctions with different ending rules and value assumptions. *Electronic Commerce Research and Applications* 14 (2), 104–111.
- Friedman, D., Rust, J., 1993. The double auction market: institutions, theories and evidence. Vol. XIV of *Santa Fe Institute Studies in the Science of Complexity*. Perseus Publishing.

- Fudenberg, D., Tirole, J., 1991. *Game Theory*. The MIT Press.
- Gerding, E. H., Dash, R., Bye, A., Jennings, N. R., 2008. Optimal strategies for bidding agents participating in simultaneous Vickrey auctions with perfect substitutes. *Journal of Artificial Intelligence Research* 32, 939–982.
- Gjerstad, S., Dickhaut, J., 1998. Price formation in double auctions. *Games and Economic Behavior* 22, 1–29.
- Gode, D. K., Sunder, S., 1993. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy* 101(1), 119–137.
- Goeree, J. K., Lien, Y., 2014. An equilibrium analysis of the simultaneous ascending auction. *Journal of Economic Theory* 153, 506–533.
- Govindan, S., Wilson, R., 2003. A global newton method to compute nash equilibria. *Journal of Economic Theory* 110 (1), 65–86.
- He, M., Jennings, N. R., Prügel-Bennett, A., 2006. A heuristic bidding strategy for buying multiple goods in multiple english auctions. *ACM Transactions on Internet Technology* 6 (4), 465–496.
- Jackson, M. O., Swinkels, J. M., 2005. Existence of equilibrium in single and double private value auctions. *Econometrica* 73 (1), 93–139.
- Lemke, C. E., Howson, J. T., 1964. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics* 12 (2), 413–423.
- Lin, R. J., Chou, S. C. T., 2004. Mediating a bilateral multi-issue negotiation. *Electronic Commerce Research and Applications* 3 (2), 126–138.
- Ma, H., Leung, H. F., 2007. An adaptive attitude bidding strategy for agents in continuous double auctions. *Electronic Commerce Research and Applications* 6 (4), 383–398.
- Mayer, B. A., Sodomka, E., Greenwald, A., Wellman, M. P., 2013. Accounting for price dependencies in simultaneous sealed-bid auctions. In: *Proceedings of the 14th ACM Conference on Electronic Commerce*. pp. 679–696.

- Miller, T., Niu, J., 2012. An assessment of strategies for choosing between competitive marketplaces. *Electronic Commerce Research and Applications* 11 (1), 14–23.
- Phelps, S., McBurney, P., Parsons, S., 2010. Evolutionary mechanism design: a review. *Autonomous Agents and Multi-Agent Systems* 21 (2), 237–264.
- Preist, C., Byde, A., Bartolini, C., 2001. Economic dynamics of agents in multiple auctions. In: *Proceedings of the 5th International Conference on Autonomous Agents*. pp. 545–551.
- Rabinovich, Z., Naroditskiy, V., Gerding, E. H., Jennings, N. R., 2013. Computing pure bayesian-nash equilibria in games with finite actions and continuous types. *Artificial Intelligence* 195, 106–139.
- Radner, R., Rosenthal, R. W., 1982. Private information and pure-strategy equilibria. *Mathematics of Operations Research* 7(3), 401–409.
- Reny, P. J., Perry, M., 2006. Toward a strategic foundation for rational expectations equilibrium. *Econometrica* 74 (5), 1231–1269.
- Sandholm, T., 2010. The state of solving large incomplete-information games, and application to poker. *AI Magazine* 31 (4), 13–32.
- Shi, B., Gerding, E. H., Vytelingum, P., Jennings, N. R., 2010. A game-theoretic analysis of market selection strategies for competing double auction marketplaces. In: *Proceedings of the 9th International Conference on Autonomous Agents and Multi-Agent Systems*. pp. 857–864.
- Shi, B., Gerding, E. H., Vytelingum, P., Jennings, N. R., 2013. An equilibrium analysis of market selection strategies and fee strategies in competing double auction marketplaces. *Autonomous Agents and Multi-Agent Systems* 26 (2), 245 – 287.
- Smith, V.L., 1962. An experimental study of competitive market behavior. *Journal of Political Economy* 70, 111–137.
- Stone, P., Schapire, R. E., Littman, M. L., Csirik, J. A., McAllester, D., 2003. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research* 19 (1), 209–242.

- Vijay, K., Sjöström, T., 1998. On the convergence of fictitious play. *Mathematics of Operations Research* 23 (2), 479–511.
- von Neumann, J., Brown, G. W., 1950. Solutions of games by differential equations. *Contributions to the Theory of Games* 24, 73–79.
- Vytelingum, P., Cliff, D., Jennings, N. R., 2008. Strategic bidding in continuous double auctions. *Artificial Intelligence* 172(14), 1700–1729.
- Wellman, M. P., Osepayshvili, A., MacKie-Mason, J. K., Reeves, D., 2008. Bidding strategies for simultaneous ascending auctions. *The B.E. Journal of Theoretical Economics* 8 (1), 27.
- Wellman, M. P., Sodomka, E., Greenwald, A., 2012. Self-confirming price prediction strategies for simultaneous one-shot auctions. In: *Proceedings of the 28th Conference on Uncertainty in Artificial Intelligence*. pp. 893–902.
- Williams, S. R., 1991. Existence and convergence of equilibria in the buyer's bid double auction. *The Review of Economic Studies* 58 (2), 351–374.
- Yuan, Y., Zeng, D., 2012. Co-evolution-based mechanism design for sponsored search advertising. *Electronic Commerce Research and Applications* 11 (6), 537–547.
- Zeng, D. D., Cox, J. C., Dror, M., 2004. Coordination of purchasing and bidding activities across markets. In: *Proceedings of the 37th Annual Hawaii International Conference on System Sciences*.