UNIVERSITY OF SOUTHAMPTON

Analysis and Design of Competing Double Auction Marketplaces

by

Bing Shi

A thesis submitted in partial fulfillment for the
degree of Doctor of Philosophy

in the
Faculty of Physical and Applied Sciences
School of Electronics and Computer Science
Intelligence, Agents, Multimedia Group

January 2011

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
SCHOOL OF ELECTRONICS AND COMPUTER SCIENCE

Doctor of Philosophy

by Bing Shi

The double auction, a highly efficient market mechanism, has been widely used by both traditional and online exchanges. However, with the globalisation of the economy, these marketplaces increasingly need to compete with each other to attract traders and charge suitable fees to make profits. In this situation, a double auction marketplace needs effective market rules (also called market policies) to govern the trading activity of its buyers and sellers and the ability to set fees appropriately in order to make profits and, at the same time, keep existing traders and attract new ones. To this end, in this thesis, we analyse competing double auction marketplaces, and use insights from this analysis to design an effective competing marketplace for an international market design competition, which is called CAT.

In more detail, the design of a competing double auction marketplace consists of determining *market policies*, which govern traders' interactions in the marketplace, and a *charging strategy*, which determines the fees charged to traders. In this thesis, we mainly focus on the latter since this is a significant determinant of the traders' choices of marketplaces and the marketplaces' profits. Now, the effectiveness of a certain charging strategy depends on the traders' behaviour, both in terms of how the fees affect their market selection, as well as their bidding behaviour. Thus, in order to set an appropriate charging strategy, we need to obtain a fundamental understanding of the traders' market selection and bidding strategies. In the context of multiple competing marketplaces, the optimal choice for a trader in terms of selecting a marketplace and submitting bids not only depends on its own preferences (i.e. type, which is usually privately known), but also on the behaviour of other traders and marketplaces, and the optimal choice of a marketplace in terms of setting fees also depends on the behaviour of traders and other marketplaces. Therefore we need to analyse the equilibrium strategies for traders and marketplaces. In so doing, we consider several settings. In particular, we consider the settings where traders can only enter one marketplace at a time (single-home trading) and can enter multiple marketplaces at a time (multi-home trading). Furthermore, we consider the setting where the traded goods are independent, substitutes or complements. In the analysis, we show how these different trading environments and different good properties affect the strategies of traders.

In more detail still, we first analyse a single-home trading environment with a small number of discrete trader types, where traders are assumed to use a truth-telling bidding strategy, i.e. submit their types as their shouts. For this setting, we first analyse the equilibrium market selection

strategies of traders for given market fees. We derive the equilibrium strategies analytically and furthermore use evolutionary game theory to investigate the dynamics of the traders' strategies. Our results show that when the same type of fees are charged by two marketplaces, all the traders will converge to one marketplace. However, when different types of fees are allowed (registration fees and profit fees), competing marketplaces are more likely to co-exist in equilibrium. Moreover, we find an interesting phenomena that sometimes all the traders eventually migrate to the marketplace that charges higher fees. We then go on to analyse the equilibrium charging strategies of the marketplaces. Specifically, we present two approaches: a static and a dynamic analysis. The former is based on the assumption that marketplaces set their fees once at the beginning and so the charging strategies are not affected by the changes in the traders' market selection strategies. In the latter analysis, we tackle this limitation by using a co-evolutionary approach where we analyse how competing marketplaces dynamically set fees while taking into account the dynamics of the traders' market selection strategies. From this analysis, we find that two initially identical marketplaces eventually charge the minimal fee that guarantees positive market profits for them. We also find an initially disadvantaged marketplace with an adaptive charging strategy can beat an initially advantaged one with a fixed charging strategy.

Building on this, we use fictitious play (a computational learning approach) to extend the above analysis by considering continuous trader types, different trading environments and different good properties. Moreover, we consider two more types of fees (transaction and transaction price percentage fees), and instead of assuming that traders adopt a truth-telling bidding strategy, we analyse both the equilibrium market selection and bidding strategies. In more detail, we first analyse traders' equilibrium bidding strategies in a single marketplace and investigate how these strategies are affected by the different fees. In so doing, we find that registration fees cause a bigger range of traders not to choose the marketplace; profit fees cause traders to shade a lot; transaction price percentage fees cause sellers to shade relatively less than buyers. Then we analyse how different trading environments and different good properties can affect traders' equilibrium market selection and bidding strategies. We then analyse the effects of different types of fees on obtaining market profits and keeping traders in a single marketplace environment. We find that the transaction price percentage fee is the most effective in making profits and keeping traders. Finally, we analyse how competing marketplaces set fees in equilibrium and show that the marketplace will charge high profit fees since traders can shade.

Finally, in addition to analysing the charging strategies, we also experimentally analyse how different market policies affect the performance of competing marketplaces in different environments where traders adopt different bidding strategies. Then, using the insights from analysing the equilibrium charging strategies and the market policies, we design a competing marketplace, which we entered into the 2010 CAT competition. This agent performed well and was ranked first in the second day's competition and second in the third day's competition.

Contents

List of Figures

List of Tables

Nomenclature

General

B	the number of buyers
$\mathcal{B} = \{1, ...B\}$	set of buyers
S	the number of sellers
$\mathcal{S} = \{1, ...S\}$	set of sellers
θ^b	a buyer's type
θ^s	a seller's type
M	the number of competing marketplaces
$\mathcal{M} = \{1, ...M\}$	set of competing marketplaces
r_m	registration fee of marketplace m
q_m	profit fee of marketplace m
t_m	transaction fee of marketplace m
o_m	transaction price percentage fee of marketplace m
p_m	fee structure of marketplace m
$\mathcal{P}$	set of all allowable fee structures
$\bar{P} = \langle p_1, ..., p_M \rangle \in \mathcal{P}^M$	fee system
k_m	pricing parameter of marketplace m
$\bar{K} = \langle k_1, ..., k_M \rangle$	pricing system
TP	transaction price

Chapter 3

F^b	type distribution function for buyers with support $[\underline{l}, \bar{l}]$
F^s	type distribution function for sellers with support $[\underline{c}, \bar{c}]$
$\omega_i^b : [\underline{l}, \bar{l}] \times \mathcal{M} \times \mathcal{P}^M \to [0, 1]$	mixed market selection strategy of buyer i
$\omega_i^s : [\underline{c}, \bar{c}] \times \mathcal{M} \times \mathcal{P}^M \to [0, 1]$	mixed market selection strategy of seller i
$\bar{\omega}^b(\bar{P}) = \langle \omega_1^b(\cdot, \bar{P}), ..., \omega_B^b(\cdot, \bar{P}) \rangle$	mixed market selection strategy profile of all buyers with $\bar{P}$
$\bar{\omega}^s(\bar{P}) = \langle \omega_1^s(\cdot, \bar{P}), ..., \omega_S^s(\cdot, \bar{P}) \rangle$	mixed market selection strategy profile of all sellers with $\bar{P}$
$\mu_m : \mathcal{P} \to [0, 1]$	mixed charging strategy of marketplace m
$\bar{\mu} = \langle \mu_1(\cdot), ...\mu_M(\cdot) \rangle$	mixed charging strategy profile of all marketplaces

$\tilde{U}_i^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b)$	expected utility of buyer i
$\tilde{U}_i^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^s)$	expected utility of seller i
$G_m^b(\theta^b \vert \bar{P})$	local type distribution of buyers in marketplace m given $\bar{P}$
$G_m^b(\theta^b)$	local type distribution of buyers in marketplace m over all possible fee systems
$\tilde{U}_m(\bar{\mu})$	expected utility of marketplace m
t_1^b	poor buyer type
t_2^b	rich buyer type
t_1^s	rich seller type
t_2^s	poor seller type
$\dot{\omega}^b(t_1^b, m, \bar{P})$	poor buyer's dynamics of market selection strategy
$\dot{\omega}^s(t_2^b, m, \bar{P})$	rich buyer's dynamics of market selection strategy
$\dot{\omega}^s(t_1^s, m, \bar{P})$	rich seller's dynamics of market selection strategy
$\dot{\omega}^s(t_2^s, m, \bar{P})$	poor seller's dynamics of market selection strategy
$\dot{\mu}_m(p_m)$	marketplace m's dynamics of charging strategy

Chapter 4

F^b	type distribution function for buyers with support $[0, 1]$
F^s	type distribution function for sellers with support $[0, 1]$
ε	small cost for traders entering marketplaces
T	units of goods traded
α_T^b	buyers' preference coefficient
$v^b(\alpha_T^b, T)$	a buyer's value on T units of goods
$D + 1$	the number of allowable shouts
$\ominus$	not choosing the marketplace
$\Phi = \{0, \frac{1}{D}, \frac{2}{D}, ..., \frac{D-1}{D}, 1\} \cup \{\ominus\}$	shout space
d_m^b	a buyer's bid in marketplace m
$\delta^b = \langle d_1^b, d_2^b, ..., d_M^b \rangle$	a buyer's action
d_m^s	a seller's ask in marketplace m
$\delta^s = \langle d_1^s, d_2^s, ..., d_M^s \rangle$	a seller's action
$\Delta = \Phi^M$	action space
ω_i^b	the probability of action i being chosen by a buyer
ω_i^s	the probability of action i being chosen by a seller
$\Omega^b = (\omega_1^b, \omega_2^b, ..., \omega_{\vert\Delta\vert}^b)$	action distributions of buyers (FP beliefs of buyers' actions)
$\Omega^s = (\omega_1^s, \omega_2^s, ..., \omega_{\vert\Delta\vert}^s)$	action distributions of sellers (FP beliefs of sellers' actions)
$\tilde{U}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P})$	expected utility of a buyer in the action distributions Ω^b and Ω^s
$\tilde{U}_m(p_m, \Omega^b, \Omega^s)$	expected utility of marketplace m

$\tilde{Q}_m(\Omega^b, \Omega^s)$	expected number of traders entering marketplace m
$\sigma^{b*}(\theta^b, \Omega^b, \Omega^s)$	a buyer's best response action against action distributions Ω^b and Ω^s
$\tilde{U}^*(\theta^b, \Omega^b, \Omega^s)$	the optimal expected utility of the buyer with type θ^b in the action distributions Ω^b and Ω^s
Ψ_i^b	type interval corresponding to best response action δ_i^b
Ω_{br}^b	best response action distribution of buyers
Ω_{br}^s	best response action distribution of sellers

Declaration of Authorship

I, Bing Shi, declare that the thesis entitled *Analysis and Design of Competing Double Auction Marketplaces* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published in a number of conference and journal papers (see Section 1.3 for a list).

Signed:

Date:

Acknowledgements

Firstly and foremost, I would like to express my sincere appreciation and gratitude to my supervisors, Professor Nicholas R. Jennings and Dr. Enrico H. Gerding, to whom I am greatly indebted for their invaluable supervision, guidance, encouragement and expertise through my PhD. They invested a tremendous amount of their time on my supervision, both during our weekly meetings and while correcting my papers and reports (often over weekends and holidays). This thesis would not have been possible without their continuous encouragement and patience.

I would also like to thank Dr. Perukrishnen Vytelingum and Dr. Ioannis Vetsikas. We cooperated to design market policies for recent years' CAT competition. They also provided me lots of valuable comments on my work. I also want to thank all the members of the IAM group who have offered their help for my research.

Furthermore, I gratefully acknowledge the financial support from Overseas Research Studentship and studentship form School of Electronics and Computer Science, University of Southampton, for funding my PhD over the past three years. This meant that I was not hampered by financial constraints and instead could dedicate my time to research.

Finally, but not least, I would like to thank my family who have provided me with unstinting support, love and care during the course of my thesis.

To my family

Chapter 1

Introduction

Exchanges, which are organised marketplaces where securities, futures, stocks and commodities can be traded, are becoming ever more prevalent. Well known traditional exchanges include the National Association of Securities Dealers Automated Quotation System (NASDAQ, a stock exchange), Chicago Mercantile Exchange (CME, a commodity exchange) and Minneapolis Grain Exchange (MGEX, a futures exchange). Now, with the development of information technology, there also exist exchanges to trade goods online. For example, Google offers DoubleClick Ad Exchange (`http://www.doubleclick.com`), which is a real-time exchange enabling large online ad publishers and ad networks and agencies to trade advertising space. Another example is FastParts (`http://www.fastparts.com`), which provides an online exchange to trade excess electronic components and used manufacturing equipment.

In the past, there was comparatively little interaction among these exchanges because of technical restrictions (such as slow information transmission and weak data processing ability). However, because of the globalised economy, now these exchanges do not exist in isolation. For example, in China, companies can be listed on both the Shanghai and the Shenzhen Stock Exchanges; in the USA, companies may be listed on both the New York Stock Exchange (NYSE) and the NASDAQ, and even in non-US marketplaces like the London Stock Exchange (LSE). Furthermore, the same commodities (such as agricultural products and precious metals) can be listed on both the New York Mercantile Exchange (NYMEX) and the CME. Now, because the same stocks and commodities can be listed on multiple exchanges, these exchanges are in competition with one another to attract companies. For example, when the National Stock Exchange (NSE) opened in India, it proceeded to claim much of the trade volume from the more established India's Bombay Stock Exchange (BSE) (Shah and Thomas, 2000). Another example is that, during the global stock market crash in 1987, unfulfilled orders on the CME overflowed onto the NYSE (Miller et al., 1988). This competition also happens between online exchanges. For example, Google's DoubleClick Ad Exchange competes against other ad exchanges, such as Microsoft's AdECN (`http://www.adecn.com`) and Yahoo!'s Right Media (`http://www.rightmedia.com`) in order to attract ad publishers and ad networks and agencies. In addition, a number of alternative trading systems, often called "dark pools", or "dark liquidity", or "dark pools of liquidity", are propagating rapidly (Carrie, 2008). Well known dark pools include Barclays Capital's Liquidity Cross, Goldman Sachs's SIGMA X, Citi's Citi Match. In contrast to the bulk of trades executed in traditional exchanges, trades in dark pools are made anonymously and are executed outside of the marketplace. Thus nobody knows who has just made a transaction and so minimum information about traders is leaked. That's why it is called a dark pool. There is evidence that these new trading systems have taken much trade volume from the traditional exchanges because of this minimum information leakage (Carrie, 2008), and now they are also competing with one another to attract traders by varying their terms of trades (e.g. Liquidity Cross operates on a continuous matching basis during market hours, while SIGMA X also offers continuous matching, and in addition, it offers X-Cross in which the matching takes place at a scheduled time).

Many of these exchanges adopt the double auction market mechanism which is a particular type of two-sided marketplace with multiple buyers (one side) and multiple sellers (the other side) (Friedman and Rust, 1993). Specifically, in such a mechanism, traders can submit offers at any time in a specified trading round and they will be matched by the marketplace at a specified time. The advantages of this mechanism are that traders can enter the marketplace at any time and they can trade multiple homogeneous or heterogeneous items in one place without travelling around several marketplaces. In addition, this mechanism is highly efficient in terms of trading goods between buyers and sellers (Smith, 1962). Since such a market mechanism has been widely used by traditional and online exchanges, in this thesis, we focus on the competition between multiple double auction marketplaces.

Specifically, in the competition between double auction marketplaces, traders that want to trade goods (stocks, commodities or advertising spaces) have a choice of marketplaces in which to participate. Thus marketplaces need to design effective and efficient market rules to govern the trading process in order to attract traders. Moreover, competing marketplaces usually charge some form of fee to the traders so they can make profits, and this choice will also affect traders' choices of marketplaces. Intuitively, we can see that there exists a conflict between making profits by charging fees and attracting traders. A marketplace can make a high short-term profit by charging high fees, but will lose traders in the long term. Thus they need to set their fees appropriately to make profits, while still maintaining the number of traders at a good level. Furthermore, a trader's choice of competing marketplaces not only depends on market rules and market fees set by the competing marketplaces, but also depends on other traders' behaviour (in terms of selecting marketplaces and making offers). In more detail, buyers(sellers) will prefer the double auction marketplace with more sellers(buyers) but less buyers(sellers) since this will increase the probability of making transactions for them. Against this background, in this thesis, we will analyse how traders behave strategically in the context of competing marketplaces, and how marketplaces compete with each other effectively in terms of establishing market rules and setting fees. We will also use insights from this formal analysis to guide the design of a practical competing marketplace agent.

Given this context, in the next section (Section 1.1), we give an overview of the competing double auction marketplace design. Then, in Section 1.2, we discuss the research challenges of this work and outline our contributions in Section 1.3. Finally, we detail the structure of the thesis (Section 1.4).

1.1 Competing Marketplace Design

In this section we outline the key components of competing double auction marketplace design. In more detail, a *marketplace* can be defined as an actual or metaphorical space, where a set of rules is established, by which buyers and sellers are in contact to exchange goods or services (Begg et al., 1994). Given this, research on marketplace design is mainly concerned with how to design these rules (also called the *market policies*) to govern traders' interactions to achieve some desirable properties, like high allocative efficiency (meaning the ratio of total profits earned by all traders to the maximum possible total profits in the marketplace) or reduced fluctuation of transaction prices.

Currently, many of the world's marketplaces are auction based, where an auction is defined as a mechanism or set of rules for trading goods (Jones, 1988). Therefore, market design is often also referred to as auction design[1]. To this end, Wurman et al. (2002) investigate the design space of auctions, which are commonly used in marketplaces in general, and identify three core activities which structure the space:

- **Receive offers:** an offer is the price at which a buyer is willing to buy a good or the price at which a seller is willing to sell a good. When receiving an offer, the auction needs to verify whether it satisfies the offer accepting rules, and if so, it will admit this offer into the active set of offers.
- **Clear:** the central purpose of an auction is to *clear* the market, i.e. execute all possible transactions (determining the allocation of goods and the corresponding payments between buyers and sellers). This activity aims to leave no possible transactions among the remaining offers. It contains three sub-activities:
  - **Timing**: indicates when to execute possible transactions.
  - **Matching**: indicates how to execute possible transactions.
  - **Pricing**: indicates the corresponding transaction prices between matched buyers and sellers.
- **Reveal intermediate information:** auctions usually supply traders with information about the state of bidding during the process, in order to guide traders toward a final outcome. These status report are called *quotes*.

[1]Note that not all marketplaces are auction based, such as traditional shopping centres, eBay's half.com and Amazon Marketplace.

Thus in the design of auction mechanisms, *offer accepting*, *timing*, *matching*, *pricing* and *quote policies* need to be determined to govern the above core activities. These policies are referred to as market polices. Now, for ease of exposition, we first use an English auction as an example to introduce these policies. In the English auction, only the bid that is higher than the current standing bid (which is the highest bid at any given moment) is accepted (offer accepting policy). If no competing buyer challenges the standing bid in a given time, then the buyer with the standing bid becomes the winner (timing policy). The item is sold to the buyer with the highest bid at a price equal to the standing bid (matching and pricing policies). During the auction, all bids at any time are public to all buyers (quote policy).

In the typical double auction marketplace, market policies are also needed to govern the above three activities. In more detail, in the double auction, any seller can submit an offer (called an *ask*) at any time in a specified trading period. This ask is observed simultaneously by all buyers and sellers. Similarly, any buyer can submit an offer (called a *bid*) at any time, which is also observed by all buyers and sellers. In the following, we use *shout* as a generic term for a bid or an ask, and call the offer accepting policy the shout accepting policy. How traders submit their shouts is determined by their *bidding strategies*. When traders attempt to place their shouts, the marketplace needs to use a shout accepting policy to determine whether to admit a trader's shout or not. Then when some shouts have been admitted, the marketplace needs to use a timing policy to determine when to match bids with asks to make transactions. For a successful match, the marketplace uses a pricing policy to determine the transaction price between the matched buyer and seller. Furthermore, the marketplace also needs to generate quote information to help traders submit shouts based on the quote policy. In order to give further understanding of these market policies as they are used in double auctions, we now introduce some examples of each of these different policies: the quote-beating accepting policy which only accepts bids higher than the bid quote and asks lower than the ask quote (see section 2.3.2.2 for more details); the round clearing and equilibrium matching policy which clears the marketplace by matching the highest bids with the lowest asks, and does so when all traders have submitted their shouts (sections 2.3.2.2 and 2.3.2.2); the *k*-pricing policy which sets the transaction price of a matched buyer and seller at some point in the interval between the buyer's bid and the seller's ask (section 2.3.2.2). Specifically, considering that different bidding strategies will be used by traders in the double auction marketplace, and these strategies rely on different information provided by the marketplace, researchers often assume that the quote policy publishes all information about the state of traders' shouts to all traders in double auctions. In our research, we also adopt this assumption, and thus do not need to consider the design of a specific quote policy.

Moreover, in addition to establishing such policies, marketplaces also need to set charging policies (also referred to as charging strategies, in this thesis we use the two terms inter-changeably), which determine fees charged to traders by marketplaces which participate in them. Such fees enable the marketplaces to earn profits and are common in real life[2]. In this thesis, we consider two categories of fees which are common in the real-world marketplaces: *ex ante* fees, which

[2]As an example, eBay charges listing fees to sellers when items are listed, and both eBay and Amazon charge final fees to sellers when items are successfully sold.

are charged to traders before they make transactions (like the listing fees adopted by eBay), and *ex post* fees, which are charged to traders after they successfully make transactions (like the final fees adopted by both eBay and Amazon). In the isolated marketplace environment, a simple charging strategy is usually sufficient since traders have no other marketplaces to select. However, in the context of competing marketplaces, traders can move freely among marketplaces to find the best deal. Therefore how marketplaces set fees becomes important since this will significantly affect the traders' choices of marketplaces and the marketplaces' profits. Thus in addition to effective market policies, a smart charging strategy is also needed in the competing double auction marketplace design.

1.2 Research Challenges

As discussed above, the design of a competing double auction marketplace consists of the design of market policies and a charging strategy. Intuitively, we can see that the effectiveness of market policies and charging strategy is affected by the traders' behaviour, which is determined by their strategies. In more detail, in the context of multiple competing marketplaces, traders need to decide which marketplaces to participate in (determined by their *market selection strategies*) and how to submit their shouts in the selected marketplaces (determined by their *bidding strategies*). Thus firstly, we need to obtain a fundamental understanding of traders' strategies in terms of market selection and bidding. Based on this, we can analyse how marketplaces should set their market policies and charging strategies to make profits while still maintaining traders.

In more detail, in the context of multiple competing double auction marketplaces, traders' market selection and bidding strategies not only depend on market policies and the charging strategies, but also depend on other traders' strategies. Specifically, there exists a *positive size effect* (Ellison et al., 2004), whereby, buyers(sellers) prefer marketplaces which have a larger number of sellers(buyers) since this gives the buyers(sellers) access to more choices. Such an effect will always push all traders towards concentrating into a single marketplace. However, in addition to the positive size effect, in double auctions, buyers(sellers) also compete with each other in order to be matched with sellers(buyers). This is referred to as a *negative size effect* (Ellison et al., 2004), whereby, traders prefer marketplaces with fewer other traders on the same side. Thus this negative size effect will push traders to distribute across different marketplaces. The positive and negative size effects have contrary impacts on the traders' distribution across multiple marketplaces, and so enhances the complexity of analysing traders' strategies (especially for market selection). Furthermore, we are interested in analysing which effect has a larger impact, and whether competing marketplaces can co-exist, and the competition can be maintained, or whether the marketplaces collapse to a monopoly setting where all traders move to one marketplace. This is important since competition drives efficiency and offers more and better choices to traders. Moreover, in contrast to much existing work that makes simplifying assumptions that all traders are homogeneous with the same preferences (i.e. types), or that marketplaces have complete information about the types of traders, many realistic marketplaces have heterogeneous traders whose types are only privately known. Specifically, in such settings, trader types can be drawn from a discrete distribution (i.e. discrete trader types), or can be drawn from a continuous distribution (i.e. continuous trader types). These assumptions also increase the complexity of the analysis.

Furthermore, in the real world, there usually exist three types of trading environments. The first is *single-home trading* where both buyers and sellers can only select one double auction marketplace at a time. The second is *multi-home trading* where both buyers and sellers can participate in multiple marketplaces at a time. The last is *hybrid trading* where one side of traders can only enter one marketplace at a time (i.e. single-home trading), while the other side of traders can enter multiple marketplaces at a time (i.e. multi-home trading). Different trading environments will affect traders' strategies, and, in turn, affect the effectiveness of the market policies and the charging strategy. For example, in a single-home trading environment, traders will only participate in the most profitable marketplace, and thus marketplaces have to compete fiercely with each other to attract traders, and have to charge low fees. However, with multi-home trading, traders will participate in any marketplace that provides non-negative (or positive) profits for them, and thus marketplaces can charge higher fees to maximise their profits. In addition to the impact of the trading environments on the marketplace competition, the properties of the goods traded between buyers and sellers can also affect the competition. Specifically, when multiple goods are traded across multiple marketplaces, these goods can be either independent, substitutes or complementary. When they are independent, the trader's value for the multiple goods is additive, i.e. equal to the sum of its value on each individual good. When they are substitutes, the trader's value is subadditive, i.e. less than the sum of its value on each individual good. When the goods are complementary, the trader's value is superadditive, i.e. greater than the sum of its value on each individual good. These different properties also affect traders' strategies. As an example, when trading complementary goods, buyers may prefer to buy as many goods as they can, and thus will try to bid high in several marketplaces to make more transactions. Therefore, when analysing competing double auction marketplaces, we need to consider these different situations.

In this thesis, we consider all the above factors when analysing competing double auction marketplaces. More specifically, the research challenges of this thesis that deal with a number of issues in the competing marketplace design are as follows:

1. ***Analyse the traders' market selection strategies:*** In the competing marketplace context, traders can move freely between marketplaces to search for the most profitable one. Intuitively, we can see that how traders select marketplaces is important since this will significantly affect the competition result of marketplaces. For example, this will affect the traders' distribution across marketplaces, and, in turn, affect the market profits. Furthermore, this will also affect the marketplaces' decisions of establishing market policies and charging strategies. Thus firstly, we should obtain a fundamental understanding about traders' market selection strategies. In the double auction with multiple buyers and multiple sellers, the optimal choice for a trader in terms of selecting a marketplace not only

depends on market policies and fees charged to it, but also depends on the behaviour of other traders. Thus we need to analyse the equilibrium market selection strategies for traders. Furthermore, as discussed above, traders' market selection strategies are affected by different trading environments and different properties of trading goods, and thus we need to consider these factors in the analysis.

2. ***Analyse the traders' bidding strategies:*** After selecting marketplaces, traders need to determine the shouts which they should submit in the chosen marketplaces. How traders submit shouts is also important since it will also affect the marketplaces' decisions of setting market policies and fees. Thus in addition to the market selection strategies, we also need to analyse the bidding strategies for traders. Similar to the market selection strategies, in double auctions, the optimal choice for a trader in terms of placing a certain shout not only depends on market policies and fees charged to it, but also depends on the shouts placed by other traders. Thus we need to analyse the equilibrium bidding strategies for traders. Currently, most work on bidding strategies is restricted to single marketplaces without considering inter-marketplace competition. Actually, even for the single marketplace, the equilibrium bidding strategy for traders is a challenging problem. In addition, in the competing marketplace context, we also need to analyse traders' equilibrium bidding strategies across multiple marketplaces. Furthermore, traders' bidding strategies are affected by the different trading environments and the different properties of trading goods as well, and thus we have to incorporate these factors in the analysis.

3. ***Analyse the market policies:*** The first part of competing market design is the choice of the market policies. A considerable body of work exists on the market policy design of an isolated double auction marketplace without inter-marketplace competition (in section 2.3.2.2, we introduce this work in detail). However, we do not know which market policy will perform well when competing with other policies. Currently, there is no systematic work on analysing the performance of market policies in the competing environment. Furthermore, the effectiveness of a market policy also depends on the traders' behaviour. Therefore, we need to analyse how different market policies affect the performance of competing marketplaces by considering different behaviours for traders.

4. ***Analyse the charging strategy:*** The second part of competing market design is the choice of a charging strategy. By charging fees to traders, the marketplaces can make profits. However, as mentioned previously, there exists a conflict between attracting traders and making profits. Thus the competing marketplace has to be able to set appropriate fees to maximise its profit, while at the same time maintain the number of traders at a good level. Furthermore, during the competition, the marketplace's opponents may change their fees and traders' bidding and market selection behaviour may change as well, and thus the marketplaces should be able to adapt their fees to these changes. Moreover, in the real world, marketplaces charge different types of fees, which usually have different effects on the traders' strategies and on obtaining profits for marketplaces. For example, when an entry fee is charged to traders for joining a marketplace, some traders may not choose

the marketplace because of the potential of negative trader profits that may be caused by this fee. Thus charging an entry fee may result in good market profit for the marketplace, but cause a decrease in the number of traders entering the marketplace. Similarly, when a percentage fee is charged on transaction profit made by traders, traders will choose this marketplace since such a fee guarantees non-negative profits for them. However, the marketplace may not be able to obtain a good market profit even by charging a high percentage fee since traders can shade their shouts to hide their actual transaction profits (as we will show in Section 4.4.1). Therefore, we need to analyse what types of fees are appropriate and effective for marketplaces to obtain market profits and maintain traders.

5. ***Design and evaluate a competing marketplace:*** Once we have analysed market policies and charging strategies and obtained insights from the analysis, we will use these insights to design a practical competing marketplace agent. Furthermore, in order to test the effectiveness of our design, we will evaluate it in the CAT competition, an international benchmarking exercise in this area.

1.3 Research Contributions

In order to reach the ultimate aim of designing an effective competing double auction marketplace, firstly, we theoretically analyse the market selection and bidding strategies for traders and charging strategies for marketplaces. In our system, intuitively, we can see that how a trader selects a marketplace and submits a shout depends on other traders' decisions, as well as the market policies and market fees. Similarly, how a competing marketplace sets its fees depends on the traders' strategies and other marketplaces' fees. Thus game theory (Binmore, 1991; 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 for theoretically analysing our system. In particular, we will analyse the Nash equilibrium bidding and market selection strategies for traders and the Nash equilibrium charging strategies for marketplaces. By so doing, this is the *first* work to comprehensively analyse competing double auction marketplaces from a theoretical perspective. Furthermore, we empirically analyse the effectiveness of different market policies in the context of an international market design competition (CAT), which is part of the Trading Agent Competition (TAC) (Cai et al., 2009). Finally, we use the insights obtained from the theoretical and empirical analysis to guide the design of a practical competing double auction marketplace.

More specifically, the research contributions of this thesis are:

1. ***We use game theory to analyse the market selection strategies for traders in the setting with discrete trader types.*** This addresses research challenge 1. This is the *first* work on analysing the equilibrium market selection strategies for traders in the context of multiple double auction marketplaces, where traders have privately know types. Here we assume

that traders use a simple, truth-telling bidding strategy, and can only select one marketplace at a time (i.e. single-home trading). Given these assumptions, we first analyse the Nash equilibrium strategies for traders' market selection for given market fees. Furthermore we analyse how traders dynamically change their market selection strategies and which of the equilibria can be reached using evolutionary game theory (EGT). We show that, when the same type of fees are charged, it is unlikely that multiple competing marketplaces will co-exist at long term despite the negative size effect; all traders will simply converge to one of the marketplaces in equilibrium. However, when different types of fees are allowed, competing marketplaces can co-exist over the long term. Counter-intuitively, we find that in certain situations all traders may converge to the marketplace that charges higher fees when the marketplace initially has a larger market share. This means that the marketplace can maintain both high number of traders and high profits. We then analyse this interesting phenomenon in more detail. Specifically, we analyse and characterise in what situations traders select the marketplace that charges higher fees and what factors affect this selection.

2. ***We game-theoretically analyse the charging strategies for marketplaces in the setting with discrete trader types.*** This addresses research challenge 4. After having established the traders' equilibrium market selection strategies (i.e. contribution 1), we proceed to analyse how competing double auction marketplaces should set their fees to make profits in equilibrium in the same setting as that in contribution 1. This is also the *first* work on analysing how two competing double auction marketplaces set fees in equilibrium. In particular, we analyse the Nash equilibrium charging strategies for marketplaces using two different approaches. In the first, we calculate marketplaces' profits for each possible type of fees, and then generate the payoff table, from which we find the equilibrium fees. However, this method does not consider the fact that marketplaces' charging strategies are affected by the dynamic changes of the traders' market selection strategies. In the second approach, we address this limitation by modelling the interplay as a two-stage game, where, in the first stage, competing marketplaces set their fees, and, in the second stage, traders select a marketplace conditional on these fees. We then use a co-evolutionary approach to analyse this game. Specifically, we find that two initially identical competing marketplaces will eventually charge the minimal fee that guarantees positive market profits for them when traders initially have an equal probability of choosing each of them. We also find that by dynamically evolving the charging strategy, it is possible for the marketplace that is initially at a disadvantage to outperform its opponent, which is also able to evolve its charging strategy. Furthermore, we show that an initially disadvantaged marketplace with an adaptive charging strategy can beat the initially advantaged one with a fixed charging strategy.

3. ***We use fictitious play to analyse the market selection and bidding strategies for traders and charging strategies for marketplaces in the setting with continuous trader types.*** This addresses research challenges 1, 2 and 4. We use fictitious play, a computational learning approach, to extend the above analysis by considering continuous trader types,

different trading environments and different properties of trading goods. Here we assume that traders adopt discrete shouts. Given this setting, we first analyse the equilibrium bidding strategies for traders in single double auction marketplaces without considering inter-marketplace competition. We show empirically that the fictitious play algorithm converges to a unique pure Bayes-Nash equilibrium. We then analyse how market fees affect these strategies. By so doing, this is the *first* work that analyses what traders will bid in equilibrium in double auctions and analyses the effect of market fees on the equilibrium bidding strategies. Building on this analysis, we go on to study competing marketplace environments where we consider both the traders' bidding and the market selection strategies. Firstly, we analyse the traders' equilibrium strategies in the single-home trading environment with independent trading goods. We find that all traders that choose a marketplace eventually converge to the same one. We then extended the analysis by considering multi-home and hybrid trading environments and different good properties, which is also the *first* work on considering these factors in the analysis of competing double auction marketplaces. Finally, we analyse what types of fees are effective for marketplaces to obtain market profits and maintain traders, and what are the equilibrium charging strategies.

4. ***We empirically analyse market policies and design an effective charging strategy for the CAT competition.*** This addresses research challenges 3 and 5. Through extensive experiments, which are based on the CAT competition platform, we obtain empirical insights about how each market policy influences the performance of competing marketplaces when different bidding strategies are adopted (GD, ZIP, RE and ZI-C bidding strategies, see Section 2.3.2.1). Examples of such insights include the fact that the allocative efficiency of a marketplace is low when it adopts a continuous clearing policy with most traders using a ZI-C or RE bidding strategy, or when the marketplace adopts an equilibrium accepting policy with most traders using a GD or ZIP bidding strategy. From these insights, we design market policies used for the CAT competition. Furthermore, based on our theoretical analysis of charging strategies, we design a novel and adaptive charging strategy in which the marketplace adjusts its fees based on the relative number of transactions. Finally, we entered our competing marketplace into the 2010 CAT competition and showed that it performed very well in this open benchmarking competition.

The following peer reviewed papers have been published or submitted to support these contributions:

- *Shi, B., Gerding, E. H., Vytelingum, P. and Jennings, N. R. (2010)* An Equilibrium Analysis of Market Selection Strategies and Fee Strategies in Competing Double Auction Marketplaces. *Submitted to the Journal of Autonomous Agents and Multi-Agent Systems.*
  In this paper, we analyse the equilibrium market selection strategies for traders in the setting with discrete trader types and use a co-evolutionary approach to analyse the equilibrium charging strategies for marketplaces. This deals with contributions 1 and 2 (see

Chapter 3).

- *Shi, B., Gerding, E. H., Vytelingum, P. and Jennings, N. R. (2010)* An Equilibrium Analysis of Competing Double Auction Marketplaces using Fictitious Play. *In: 19th European Conference on Artificial Intelligence (ECAI), Lisbon, Portugal. pp. 575-580.*
  In this paper, we use fictitious play to analyse the equilibrium bidding and market selection strategies for traders in the setting with continuous trader types. This partly deals with contribution 3 (see Chapter 4).

- *Shi, B., Gerding, E. H., Vytelingum, P. and Jennings, N. R. (2010)* A Game-Theoretic Analysis of Market Selection Strategies for Competing Double Auction Marketplaces. *In: 9th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS2010), Toronto, Canada. pp. 857-864.*
  In this paper, we game theoretically analyse Nash equilibrium market selection strategies for traders in the setting with discrete trader types, and use EGT to analyse the dynamics of traders' strategies. This deals with contribution 1 (see Chapter 3).

- *Shi, B., Gerding, E. H., Vytelingum, P. and Jennings, N. R. (2010)* Setting Fees in Competing Double Auction Marketplaces: An Equilibrium Analysis. *In: 12th International Workshop on Agent-Mediated Electronic Commerce (AMEC 2010), Toronto, Canada. pp. 85-98.*
  In this paper, we game theoretically analyse Nash equilibrium charging strategies for traders given traders' equilibrium market selection strategies. This deals with contribution 2 (see Chapter 3).

- *Vytelingum, P., Vetsikas, I., Shi, B. and Jennings, N. R. (2008)* IAMwildCAT: The Winning Strategy for the TAC Market Design Competition. *In: 18th European Conference on Artificial Intelligence (ECAI), Patras, Greece. pp. 428-432.*
  In this paper, we empirically analyse the market policies used in the CAT competition context. This partly deals with contribution 4 (see Chapter 5).

1.4 Thesis Structure

The outline of the thesis is as follows:

- In Chapter 2, we provide the necessary background on game theory and markets, and analyse the literature about both isolated and competing marketplaces. Finally, we describe the CAT competition and review relevant work in this competition.

- In Chapter 3, we game theoretically analyse the equilibrium market selection strategies for traders and the equilibrium charging strategies for marketplaces in the setting with discrete trader types.

- In Chapter 4, we use fictitious play to analyse the equilibrium market selection and bidding strategies for traders in the setting with continuous trader types, different trading environments and different good properties. We also analyse the effects of different types of fees on obtaining market profits. Furthermore, we analyse how competing marketplaces set fees in equilibrium.

- In Chapter 5, we empirically analyse the performance of different market policies when different bidding strategies are adopted, and use insights from this analysis to design market policies for the CAT competition. We also use insights from the theoretical analysis of the charging strategy to guide the practical design of a novel and adaptive charging strategy. We show that our design of market policies and the charging strategy performed well in the 2010 CAT competition.

- In Chapter 6, we conclude this thesis and outline future work.

Chapter 2

Literature Review

In this chapter, we begin by introducing the basic background on game theory (Section 2.1) and general market theory from microeconomics (Section 2.2). These topics are relevant because they form the foundations for analysing marketplace competition. Then, we go on to describe related literature about isolated marketplaces, including bidding strategies and market policies (Section 2.3). Following this, we review work on competing marketplaces (Section 2.4). We then introduce the Market Design Competition (Section 2.5). Finally, we summarise this chapter (Section 2.6).

2.1 Background on Game Theory

Game theory, which studies the strategic interactions of self-interested agents mathematically, is an important part of microeconomics. Thus it has been widely used to analyse the interactions between traders and between competing marketplaces. In this section, we provide the basic notions from game theory which are related to our work. For a comprehensive overview of this area, see Fudenberg and Tirole (1991) and Osborne and Rubinstein (1994).

In more detail, a game consists of a set of players which we will denote by $\mathcal{I} = \{1, 2, ..., I\}$. Each player has a strategy, which is a complete contingent plan, or decision rule, to specify how this player will act in each possible distinguishable circumstance. Specifically, we use s_i to represent a strategy for player i, and $s_i \in S_i$, where S_i is the set of all possible strategies for player i. In the game, the strategies of all players constitute a strategy profile, which is an I-tuple, $\bar{s} = \langle s_1, ..., s_I \rangle \in S$, where $S = S_1 \times ... \times S_I$ is the set of all possible strategy profiles. In addition, we use $s_{-i} = \langle s_1, ..., s_{i-1}, s_{i+1}, ..., s_I \rangle$ to represent the strategy profile for all players except i, and $s_{-i} \in S_1 \times ... \times S_{i-1} \times S_{i+1} \times ... \times S_I$. Now $\bar{s}$ can be rewritten as $\bar{s} = \langle s_i, s_{-i} \rangle$. Each strategy profile will induce an outcome for the game. Players have different preferences over different outcomes. We use utility functions to describe players' preferences over outcomes. Formally, the player i's utility function is defined as $U_i : S \rightarrow R$, which is a mapping from the set of players' strategy profiles to the utilities over the outcomes induced by the strategy profiles. Now, given the set of players, the players' strategy profiles and the utility functions, a game can be formally represented in the following way:

Definition 2.1. For a game with player set $\mathcal{I}$, Γ specifies for each player i a set of strategies S_i (with $s_i \in S_i$) and a utility function $U_i(\langle s_1, ..., s_I \rangle)$. Formally, the game can be written as $\Gamma = [\mathcal{I}, \{S_i\}, \{U_i(\cdot)\}]$, $i \in \mathcal{I}$.

The strategy we presented above is termed a *pure* strategy, where players choose strategies in a deterministic way. However, stable outcomes (e.g. Nash equilibria, as we will introduce below) where all players play pure strategies, do not always exist. Therefore, in addition to pure strategy, we also introduce the *mixed* strategy, whereby a probability is assigned to each pure strategy. The mixed strategy allows a player to randomly select a pure strategy. Formally, a mixed strategy of player i can be defined as $\omega_i : S_i \rightarrow [0, 1]$, which assigns to each pure strategy $s_i \in S_i$ a probability $\omega_i(s_i) \geq 0$ that it will be played, where $\sum_{s_i \in S_i} \omega_i(s_i) = 1$. Here, it is important to note that a pure strategy can be regarded as a degenerate case of a mixed strategy, where the particular pure strategy is selected with probability 1 and every other strategy with probability 0. Because probabilities are continuous, there are infinitely many mixed strategies available to a player, even when the pure strategy set is finite. The set of possible mixed strategies for player i which has M pure strategies in set $S_i = \{s_{i\_1}, ..., s_{i\_M}\}$ can be represented by

$$\Delta_i = \Big\{(\omega_i(s_{i\_1}), ..., \omega_i(s_{i\_M})) \in R^M : \omega_i(s_{i\_m}) \geq 0 \; for\ all\ m \in \{1, ..., M\}\ and\ \sum_{m=1}^{M} \omega_i(s_{i\_m}) = 1\Big\}$$

Similarly, a mixed strategy profile can be denoted by $\bar{\omega} = \langle \omega_1, ...\omega_I \rangle \in \Delta$, where $\Delta = \Delta_1 \times ... \times \Delta_I$ is the set of all possible mixed strategy profiles. Furthermore, the mixed strategy profile can be rewritten as $\bar{\omega} = \langle \omega_i, \omega_{-i} \rangle$, where $\omega_{-i} \in \Delta_1 \times ... \times \Delta_{i-1} \times \Delta_{i+1} \times ... \times \Delta_I$ is the mixed strategy profile for all players except i. Now, the players' utility functions need to be redefined since their utilities are in expectation over probability distributions of pure strategies. Specifically, player i's expected utility function is defined as $\tilde{U}_i : \Delta \rightarrow R$. In more detail, it is the expectation over the probability distribution of pure strategy profile $\bar{s} = \langle s_1, ..., s_I \rangle$:

$$\tilde{U}_i(\bar{\omega}) = \sum_{\bar{s} \in S} \prod_{i \in \mathcal{I}} \omega_i(s_i) * U_i(\bar{s})$$

Therefore, the game where players adopt mixed strategies can be written as $\Gamma = [\mathcal{I}, \{\Delta_i\}, \{\tilde{U}_i(\cdot)\}]$, $i \in \mathcal{I}$.

After presenting the game with mixed strategies, we now introduce the notion of *solution concept*. In this context, a solution concept is a formal rule used to predict how the game will be played. These predictions describe which strategies will be adopted by the players, therefore predicting the outcome of the game. In the following, we describe the two most commonly used solution concepts: *dominant strategy* and *Nash equilibrium*. For the solution concept of dominant strategy, it means that the player can maximise its utility by adopting a certain strategy no matter what strategies other players use. Formally, a dominant strategy is defined as:

Definition 2.2. A pure strategy $s_i \in S_i$ is a dominant strategy for player i in game $\Gamma = [\mathcal{I}, \{S_i\}, \{U_i(\cdot)\}]$, $i \in \mathcal{I}$, if

$$\forall s_i' \in S_i \textit{ and } s_i' \neq s_i, \forall s_{-i} \in S_{-i}, U_i(\langle s_i, s_{-i}\rangle) \geq U_i(\langle s_i', s_{-i}\rangle)$$

Generally speaking, a dominant strategy is a very robust solution concept since it is not based on any assumptions about the information available to players about each other, and also does not require each player to believe that other players will adopt its own optimal strategy. However, the shortcoming is that dominant strategies often do not exist.

The other main solution concept, the Nash equilibrium, requires that in equilibrium, each player will select the strategy which maximises its (expected) utility given the strategies of all other players. Formally, a pure strategy Nash equilibrium is defined as:

Definition 2.3. A pure strategy profile $\bar{s}^* = \langle s_1^*, ..., s_I^*\rangle$ constitutes a pure strategy Nash equilibrium of game $\Gamma = [\mathcal{I}, \{S_i\}, \{U_i(\cdot)\}]$, $i \in \mathcal{I}$, if

$$\forall i \in \mathcal{I}, \forall s_i \in S_i, U_i(\langle s_i^*, s_{-i}^*\rangle) \geq U_i(\langle s_i, s_{-i}^*\rangle)$$

This definition can be extended to the game including mixed strategies:

Definition 2.4. A mixed strategy profile $\bar{\omega}^* = \langle \omega_1^*, ..., \omega_I^*\rangle$ constitutes a mixed strategy Nash equilibrium of game $\Gamma = [\mathcal{I}, \{\Delta_i\}, \{\tilde{U}_i(\cdot)\}]$, $i \in \mathcal{I}$, if

$$\forall i \in \mathcal{I}, \forall \omega_i \in \Delta_i, \tilde{U}_i(\langle \omega_i^*, \omega_{-i}^*\rangle) \geq \tilde{U}_i(\langle \omega_i, \omega_{-i}^*\rangle)$$

Furthermore, in some games, it may be too complicated to derive the Nash equilibrium. Therefore researchers have to approximate the Nash equilibrium. Specifically, ϵ-Nash equilibrium is proposed as a solution concept which approximately satisfies the condition of Nash equilibrium (Leyton-Brown and Shoh, 2008). It is formally defined as:

Definition 2.5. A pure strategy profile $\bar{s}^* = \langle s_1^*, ..., s_I^*\rangle$ constitutes a pure strategy ϵ-Nash equilibrium of game $\Gamma = [\mathcal{I}, \{S_i\}, \{U_i(\cdot)\}]$, $i \in \mathcal{I}$, if

$$\forall i \in \mathcal{I}, \forall s_i \in S_i, U_i(\langle s_i^*, s_{-i}^*\rangle) \geq U_i(\langle s_i, s_{-i}^*\rangle) - \epsilon$$

It means that in the ϵ-Nash equilibrium, it is not possible for any player to gain more than ϵ in expected utility by unilaterally deviating from its strategy. This definition for the game with mixed strategies is analogous. Note that every Nash Equilibrium is equivalent to a ϵ-Nash equilibrium where $\epsilon = 0$. In Chapter 4, we use this concept to approximate the Nash equilibrium strategies for traders.

So far, in the description of Nash equilibrium, we have assumed that each player knows the relevant information of all players, including their preferences on the outcomes (i.e. utility functions). Such games are known as games of *complete information*. However, this is a very strong assumption that does not hold in all settings. For example, in a double auction marketplace, a trader may not know other traders' preferences for the goods. To overcome this shortcoming, the extension of Nash equilibrium, Bayes-Nash equilibrium was introduced by Harsanyi (1962). Formally, in a Bayesian game, each player is assumed to share common knowledge about the probability distributions of players' preferences (which are also formally referred to as types) $F(\theta_1, ..., \theta_I)$, where θ_i denotes player i's type. $\theta_i \in \Theta_i$, where Θ_i is the set of all possible types for player i. Furthermore, in a Bayesian game, a pure strategy for player i is a function $s_i(x)$, which specifies the player's action choice for each possible type $x \in \Theta_i$. The set of all possible pure strategies for player i is denoted by S_i. Moreover, in a Bayesian game, although a player does not know other player's types, it usually assumes that each player knows its own type. The expected utility of a player i with type θ_i given a strategy profile, $\bar{s} = \langle s_1(\cdot), ..., s_I(\cdot)\rangle$, is then described as:

$$\tilde{U}_i(\langle s_1(\cdot), ..., s_I(\cdot)\rangle, \theta_i) = E_{\theta_{-i}}[U_i(\langle s_1(\theta_1), ...s_I(\theta_I)\rangle, \theta_i)]$$

which is in expectation over probability distributions of all other players' types . $U_i(\langle s_1(\theta_1), ..., s_I(\theta_I)\rangle, \theta_i)$ is the player i's utility on the outcome induced by players' strategy profile $\langle s_1(\cdot), ..., s_I(\cdot)\rangle$ on the specific realisation of players' types $\theta_1, ..., \theta_I$. Then a Bayesian game can be defined as $[I, \{S_i\}, \{\tilde{U}_i(\cdot)\}, \Theta, F(\cdot)]$, $i \in \mathcal{I}$, where $\Theta = \Theta_1 \times ... \times \Theta_I$.

Furthermore, the mixed strategy for player i in the Bayesian game can be defined as a function $\omega_i(y)$, which specifies the probability of the pure strategy $y \in S_i$ being used by player i. $\omega_i(y) \in \Delta_i$, where Δ_i is the set of all possible mixed strategies for player i. Then the expected utility of a player i with type θ_i given a mixed strategy profile $\langle \omega_1(\cdot), ..., \omega_I(\cdot)\rangle$ can be described as:

$$\tilde{U}_i(\langle \omega_1(\cdot), ..., \omega_I(\cdot)\rangle, \theta_i) = \sum_{\bar{s} \in S} \omega(\bar{s}) * E_{\theta_{-i}}[U_i(\langle s_1(\theta_1), ...s_I(\theta_I)\rangle, \theta_i)]$$

which is in expectation over probability distributions of pure strategy profiles and over probability distributions of other players' types. $\omega(\bar{s}) = \prod_{i \in \mathcal{I}} \omega_i(s_i)$ is the probability of the pure strategy profile $\bar{s}$ taking place. Then a Bayesian game with mixed strategies can be defined as $[I, \{\Delta_i\}, \{\tilde{U}_i(\cdot)\}, \Theta, F(\cdot)]$, $i \in \mathcal{I}$, where $\Theta = \Theta_1 \times ... \times \Theta_I$.

In the Bayes-Nash equilibrium, each player selects a strategy to maximise its expected utility given the expected utility maximising strategies of other players. Formally:

Definition 2.6. A pure strategy profile $\langle s_1^*(\cdot), ..., s_I^*(\cdot)\rangle$ constitutes a pure strategy Bayes-Nash equilibrium of Bayesian game $\Gamma = [I, \{S_i\}, \{\tilde{U}_i(\cdot)\}, \Theta, F(\cdot)]$, if

$$\forall i \in \mathcal{I}, \forall \theta_i \in \Theta_i, \forall s_i(\cdot) \in S_i, \tilde{U}_i(\langle s_i^*(\cdot), s_{-i}^*(\cdot)\rangle, \theta_i) \geq \tilde{U}_i(\langle s_i(\cdot), s_{-i}^*(\cdot)\rangle, \theta_i)$$

Similarly, the definition can be extended to the game with mixed strategy:

Definition 2.7. A mixed strategy profile $\langle \omega_1^*(\cdot), ..., \omega_I^*(\cdot)\rangle$ constitutes a mixed strategy Bayes-

Nash equilibrium of Bayesian game $\Gamma = [\mathcal{I}, \{\Delta_i\}, \{\tilde{U}_i(\cdot)\}, \Theta, F(\cdot)]$, if

$$\forall i \in \mathcal{I}, \forall \theta_i \in \Theta_i, \forall \omega_i(\cdot) \in \Delta_i, \tilde{U}_i(\langle \omega_i^*(\cdot), \omega_{-i}^*(\cdot) \rangle, \theta_i) \geq \tilde{U}_i(\langle \omega_i(\cdot), \omega_{-i}^*(\cdot) \rangle, \theta_i)$$

Furthermore, in some games, there may exist multiple Nash equilibria. For this situation, researchers usually assume that players with the same type will adopt the same strategy in equilibrium, which is referred to as *symmetric Nash equilibrium*. In our work on analysing competing marketplaces with incomplete information, as is common in game theory, we will also focus on the symmetric Nash equilibrium of traders' market selection and bidding strategies (see Chapters 3 and 4).

2.1.1 Evolutionary Game Theory

Up to this point, we have introduced the key relevant definitions of game theory, which will be useful in our analysis of competing marketplaces. However, the Nash equilibrium makes an assumption that all players have the abilities to correctly anticipate the opposing players in equilibrium. It fails to reflect the fact that in the real world, people may not make decisions under this assumption. Furthermore, traditional game theory only provides a static explanation for why populations playing Nash equilibrium strategies remain in that state since each population makes a best response to the other populations' strategies. It fails to indicate whether the Nash equilibrium strategies can be reached and which of these equilibria is most likely to occur. To address these limitations, Maynard (1982) adopted the idea of evolution from biology to game theory, which originated evolutionary game theory (EGT). In the biological circumstances, it is also impossible to determine what decisions are the most rational ones, and therefore each "player" has to learn to optimise its strategy and maximise its utility (i.e evolution). Furthermore, the basic techniques developed in EGT were also initially formulated in the context of evolutionary biology (Maynard, 1982; Weibull, 1996). Specifically, in EGT, a dynamic process is constructed where the proportions of various strategies in a population evolve. In more detail, *replicator dynamics* are adopted in EGT to specify how players gradually adjust their strategies over time in response to the repeated observation of their opponents' strategies. They are formalised as a system of differential equations. In the following, we will introduce the replicator dynamics equation for a single and a multiple population setting.

In a single population setting, players have the same set of pure strategies to employ. Furthermore, since a very large population is being considered, the population proportion of using different pure strategies is equivalent to the mixed strategy[1]. Specifically, players are assumed to choose a strategy from the set of M pure strategies $\{s_1, ..., s_M\}$, and the probability of the strategy s_m being used is represented by ω_m. Then the mixed strategy can be denoted as $\omega = (\omega_1, ..., \omega_M)$, where $\sum_{m=1}^{M} \omega_m = 1$. The replicator dynamics will specify the dynamic adjustment of the prob-

[1]Note that the actual number of players in the game is not the size of the population. However, in the evolution, players' strategies are sampled from this large population.

ability of which pure strategy should be played, and it is defined as $\dot{\omega} = (\dot{\omega}_1, ..., \dot{\omega}_M)$, which describes how players gradually evolve their strategies over time. In more detail, the element $\dot{\omega}_m$ in $\dot{\omega}$ is defined as follows:

$$\dot{\omega}_m = \frac{d\omega_m}{dt} = [\tilde{U}(s_m, \omega) - \tilde{U}(\omega, \omega)] * \omega_m \qquad (m = 1, ..., M) \tag{2.1}$$

where $\tilde{U}(s_m, \omega)$ is the expected utility of a player using pure strategy s_m when all other players employ the mixed strategy ω, and $\tilde{U}(\omega, \omega)$ is the expected utility for the mixed strategy ω. To get the dynamics of the game, one needs to calculate *trajectories*, which indicate how the mixed strategies evolve. In more detail, initially, a mixed strategy ω is randomly chosen as a starting point. The dynamics $\dot{\omega}$ is then calculated according to equation 2.1. According to the changes of the probabilities of which pure strategy should be played, their current mixed strategy can be updated. This calculation is repeated until $\dot{\omega}$ becomes a zero vector. At this moment, the equilibrium is reached, which is the current mixed strategy. The replicator dynamics show the trajectories and how they converge to a Nash equilibrium. We call a Nash equilibrium to which trajectories converge, an *attractor*, and call a Nash equilibrium to which no trajectories converge, a *saddle point*. The region where all trajectories converge to a particular equilibrium is called the *basin of attraction* of this equilibrium. The basin is very useful since when each starting point is selected by players with an equal probability, its size indicates how likely the population is to converge to that equilibrium (Bullock, 1997).

In the above, we introduced replicator dynamics where interaction takes place between players from the same population, in which players evolve their mixed strategies in the same way. However, in many games, interaction may take place between players from different populations, in which players from different populations evolve their mixed strategies differently. To this end, we introduce replicator dynamics equations in the setting with two different populations. Specifically, players from the two populations are assumed to choose strategies from $\{s_1, ..., s_M\}$ and $\{s'_1, ..., s'_{M'}\}$ respectively. $\omega = (\omega_1, ..., \omega_M)$ and $\omega' = (\omega'_1, ..., \omega'_{M'})$ are used to denoted the mixed strategies used by players from the two populations respectively. Then the replicator dynamics of the two populations are $\dot{\omega} = (\dot{\omega}_1, ..., \dot{\omega}_M)$ and $\dot{\omega}' = (\dot{\omega}'_1, ..., \dot{\omega}'_{M'})$ respectively. Their individual elements $\dot{\omega}_m$ and $\dot{\omega}'_{m'}$ are calculated as follows:

$$\dot{\omega}_m = \frac{d\omega_m}{dt} = [\tilde{U}(s_m, \omega, \omega') - \tilde{U}(\omega, \omega, \omega')] * \omega_m \qquad (m = 1, ..., M) \tag{2.2}$$

$$\dot{\omega}'_{m'} = \frac{d\omega'_{m'}}{dt} = [\tilde{U}'(s'_m, \omega, \omega') - \tilde{U}'(\omega', \omega, \omega')] * \omega'_{m'} \qquad (m' = 1, ..., M') \tag{2.3}$$

where $\tilde{U}(s_m, \omega, \omega')$ is the expected utility of a player from a population using pure strategy s_m when all other players from this population use the mixed strategy ω and all players from the other population use the mixed strategy ω', and $\tilde{U}(\omega, \omega, \omega')$ is the expected utility for the player using the mixed strategy ω. Using equations 2.2 and 2.3, we can then calculate the dynamics of the two populations' strategies starting from various initial mixed strategies. From these equations, we can see that the dynamic change of a mixed strategy in each population is determined by the mixed strategy of the other population. Furthermore, the replicator dynamics equations for the two population setting can easily be extended in the same way to more population settings.

Evolutionary game theory has been widely used to study the interaction of self-interested agents (e.g. see Tuyls 2004; Phelps 2008). In Chapter 3, we will use evolutionary game theory to analyse the dynamics of traders' market selection strategies and find which equilibrium is the most likely to occur. Furthermore, we use will evolutionary game theory to analyse the marketplaces' equilibrium charging strategies from a static and a dynamic approach respectively.

2.1.2 Fictitious Play

In the above, we have introduced traditional game theory and evolutionary game theory. However, finding a Nash equilibrium is usually computationally demanding. The Lemke-Howson algorithm (Lemke and Howson, 1964), perhaps the best-known algorithm for computing Nash equilibrium, has been shown to require exponential time on some instances (Savani and Stengel, 2006). Given this, one possible way is to approximate the Nash equilibrium, i.e. compute the ϵ-Nash equilibrium. There exist a number of algorithms to approximate Nash equilibrium (Spirakis, 2008). A well-known algorithm for approximating Nash equilibrium is fictitious play, which is a computational learning approach (von Neumann and Brown, 1950; Brown, 1951). In the following, we describe the basics of this algorithm in detail.

In the standard FP algorithm (von Neumann and Brown, 1950; Brown, 1951), opponents are assumed to play a mixed strategy. Then by observing relative appearance 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. At this moment, all players' best response strategies constitute a pure Nash equilibrium. Second, it may converge in FP beliefs. At this moment, the converged FP beliefs constitute a mixed Nash equilibrium. However, 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 limited number of rounds.

However, we should note that the standard FP algorithm is not suitable for analysing Bayesian games in which there is incomplete information (i.e. the player's type is not known to the other players). 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. (2009) 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 Nash equilibrium always exists given the conditions that the game is non-atomic, giving zero probability to any specific player type to appear, and the action space is finite. However, in Rabinovich et al. (2009), researchers only showed how to use this algorithm to analyse traders' strategies in single-sided auctions. Building on this, in Chapter 4, we will describe how to apply this fictitious play algorithm to approximate the Bayes-Nash equilibrium market selection and bidding strategies for traders in the much more complex environment with multiple competing double auction marketplaces.

2.2 Background on Markets

After introducing the key basic definitions of game theory, we now give an overview on general market theory from microeconomics. The key concepts introduced in the general market theory will help us understand the double auction marketplace. In such settings, a market consists of one or more buyers and sellers. Typically, each seller has a *cost price* for the item it possesses. This is the lowest price it is willing to sell the item for. Similarly, each buyer has a *limit price* for the item it wishes to buy, which is the highest price that it is willing to buy the item for. The cost(limit) prices of traders are denoted by *types* in game theory. At each possible price, the quantity of a commodity buyers wish to buy is referred to as *demand*, and the quantity of a commodity that sellers want to sell is referred to as *supply*. Usually, the greater the price of the commodity, the lower the demand, and the lower the price, the higher the demand. Such characteristics of demand and supply are often represented by an *underlying demand and supply curve*[2], which is a function of demand and supply with respect to price (see Figure 2.1).

Given a stable underlying demand and supply, classical microeconomic theory claims that in a marketplace with profit-motivated[3] traders, transaction prices will converge to an *equilibrium price* p^*, where the quantity of demand is equal to the quantity of supply (Mas-Collel et al., 1995). This quantity is referred to as the *equilibrium quantity* q^*. The reason why transaction prices are beheld to converge to an equilibrium price is as follows. In a marketplace with profit-motivated traders, when transaction prices are below the equilibrium price, there is excess demand. Because of this, sellers can raise their asks, and then in order to remain competitive, buyers have an incentive to bid higher. This will raise the transaction prices towards the equilibrium price. Similarly, when prices are above the equilibrium price, there is excess supply and buyers can buy items from sellers with lower prices. This means that sellers have an incentive to reduce their asks to compete against other sellers, thus lowering transaction prices towards the equilibrium price. At the equilibrium price, neither buyers nor sellers have any incentive to

[2]Here, 'underlying' means that the demand and supply curve is determined by the buyers' limit prices and the sellers' cost prices. It may be different from the *reported demand and supply curve* which is determined by traders' shouts.

[3] That is, a trader will always choose the action that maximises its expected profit.

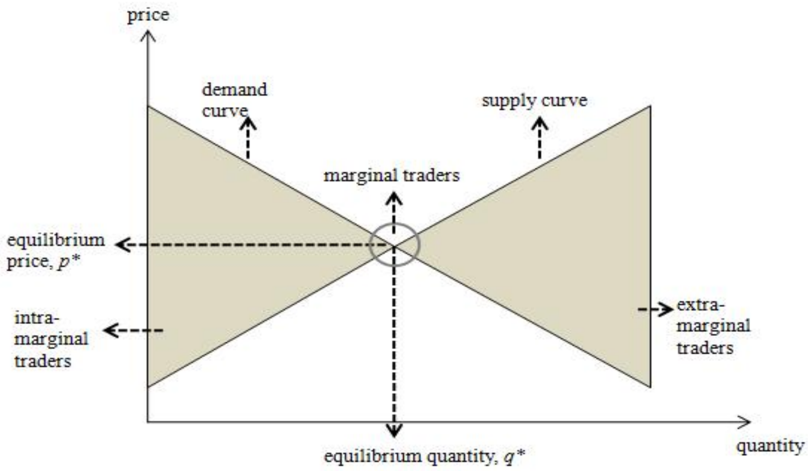

FIGURE 2.1: Demand and supply curve.

change their prices, and so the marketplace becomes stable. The value of the equilibrium price and quantity can be determined according to the intersection of the demand and supply curve (see Figure 2.1). In this context, we distinguish between two types of traders: *intra-marginal* and *extra-marginal*. Traders to the left of the equilibrium point are known as *intra-marginal buyers* (whose limit prices are higher than p^*) and *intra-marginal sellers* (whose cost prices are less than p^*). On the right of the equilibrium point are extra-marginal buyers (whose limit prices are less than p^*) and extra-marginal sellers (whose cost prices are higher than p^*). Intra-marginal traders are more likely to make transactions than extra-marginal traders in practice. In addition, we often refer to traders around the equilibrium point as marginal traders. These traders are less likely to make transactions than intra-marginal traders, but more likely to make transactions than extra-marginal traders.

In the marketplace, when an equilibrium quantity of goods is traded at the equilibrium price, the optimal allocation is reached, which means that the *allocative efficiency* is maximised. Specifically, the allocative efficiency E is the total profit earned by all traders in the marketplaces divided by the maximum possible total profit that could have been earned by all traders:

$$E = \frac{\sum_{i \in T} |v_i - \mathrm{TP}_i|}{\sum_{i \in T^*} |v_i - p^*|} \times 100\% \tag{2.4}$$

where T is the set of traders making transactions in the marketplace, v_i is the cost or limit price of trader i, TP_i is the transaction price of trader i, $|v_i - \mathrm{TP}_i|$ is trader i's actual profit in the transaction, T^* is the set of intra-marginal traders, $|v_i - p^*|$ is i's expected profit if it trades at the equilibrium price. In a marketplace where all buyers' limit prices and all sellers' cost prices are public, i.e. the underlying demand and supply are known to all traders and the marketplace, the equilibrium quantity and equilibrium price can be computed, and an optimal resource allocation can be performed.

2.3 Isolated Auction Marketplaces

Thus far we have introduced the demand and supply model of marketplaces, and described the concept of equilibrium price and equilibrium quantity. Now, before we review related work about competing marketplaces, we first introduce related work about isolated marketplaces without inter-marketplace competition since the understanding of isolated marketplaces is the foundation for further understanding competing marketplaces. Furthermore, in this thesis, since we focus on the auction-based marketplaces, here we introduce relevant work on isolated auctions. Note that, although in this thesis we focus on the double auction marketplaces, the work in the single-sided auction marketplaces will be important for us to understand the double auction marketplaces. Therefore, here we first introduce related work about isolated single-sided auctions, and then introduce isolated double auction marketplaces.

2.3.1 Single-Sided Auctions

In a single-sided auction, there is either a single seller and multiple buyers competing to win the good, or there are multiple sellers competing to provide a good or service to a single buyer (also called a reverse or procurement auction). In the following, we will describe the formats of some widely used single-sided auctions. Specifically, we introduce the auctions where each buyer has a *private value*[4] on the good, i.e. each buyer knows the value of the good to himself at the time of bidding, but he does not know the exact values of other buyers to the good, and knowledge of other buyers' values will not affect his own value.

The open ascending price or English auction is probably the oldest and most prevalent auction form in the world (Krishna, 2002). In the traditional English auction, only the bid that is higher than the current standing bid (which is the highest bid at any given moment) is accepted by the auctioneer. In a given time, if no competing buyer challenges the standing bid, then the buyer with the standing bid will win the good and pay a price equal to the standing bid. Furthermore, in addition to the traditional English auction, there exist many variations on this auction system. For example, in eBay, there is a deadline for selling the good through the auction. In a Japanese auction, each buyer should indicate his interest in purchasing the good at the current price. As the price rises, he may indicate that he is no longer interested and quits the auction. Once he quits, he cannot reenter the auction. The auction continues until only one buyer remains. The English auction is commonly used for selling antiques and artwork, but also for selling used goods and real estate. Another well-known auction is the Dutch auction, which has a descending price. In this auction, the auctioneer begins by calling out a price high enough so that initially no buyer is interested in buying the good at that price. This price is then gradually lowered until some buyer indicates his interest. The good is then sold to this buyer at the given price. This type of auction is usually used to sell perishable commodities, such as fish, flowers and tobacco.

[4]This is the limit price for the buyer, and is also referred to as the type of the buyer, see Section 2.2.

Note that both English and Dutch auctions are open auctions where in the auction process, each buyer's behaviour is observed by all other buyers and the auctioneer. There also exist two well-known sealed-bid auctions: first-price and second-price. In both auctions, buyers submit bids in sealed envelopes and the buyer submitting the highest bid wins the good. In contrast to the English auction, in these auctions, buyers can only submit one bid. Furthermore, as buyers cannot see the bids of other participants, they cannot adjust their own bids accordingly. Now, in the first-price sealed-bid auction, the winner will pay what he actually bids. In the second-price sealed-bid auction (also called Vickrey auction), the buyer submitting the highest bid will win the good, but pays the second highest bid. Although this auction is extremely important in auction theory, in the real world, Vickrey auctions are rarely used.

After describing the above four common single-sided auction mechanisms, we now describe the bidding strategies used in these auctions. Krishna (2002) points out that the Dutch auction is strategically equivalent to the sealed-bid first-price auction. The reason is as follows. Although in the Dutch auction, each buyer can observe that some buyer has agreed to buy at the current price, since that causes the auction to end, the buyer cannot utilise such information to improve his bid. Thus bidding a certain amount in a sealed-bid first-price auction is equivalent to bidding to buy at that amount in a Dutch auction. Furthermore, when buyers have private values on the goods, the English auction is also weakly equivalent to the sealed-bid second-price auction. Because of the strategic equivalence of the different single-sided auctions, in the following, we introduce equilibrium bidding strategies in the two sealed-bid auctions.

In the auctions, a bidding strategy for a buyer is defined as a function mapping the set of possible private values (types) to a set of allowable bids. In more detail, in a sealed-bid second-price auction, the dominant strategy is to bid buyers' private values truthfully. In a sealed-bid first-price auction, the equilibrium behaviour is more complicated than in a second-price auction. Intuitively, no buyers will bid their private value truthfully since this causes zero utility if he wins. Therefore, in the equilibrium, buyers should shade their bids (bid is less than the actual private value), and the degree of shading depends on the number of competing buyers and the probability distributions of buyers' private values. When the number of buyers in the auction increases, the degree of shading approaches 0, i.e. in this situation, buyers' bids approach to their true private values.

In the above, we have described single-sided auctions with private values. There also exist settings with a common value, where the value, although unknown at the time of bidding, is the same for all buyers. Furthermore, in addition to a single good traded in the auction, there exist variations in which multiple goods can be traded. However, it is beyond the scope of this review to discuss these variations. Details about these auctions can be found in Krishna (2002).

2.3.2 Double Auctions

In the above, we have introduced single-sided auctions. Now we introduce related work about double auction marketplaces. As we discussed in Chapter 1, the double auction market mechanism is widely used in stock exchanges, commodity exchanges and even online exchanges. Specifically, a double auction marketplace is a particular type of two-sided marketplace with multiple buyers (one side) and multiple sellers (the other side). In such a marketplace, both buyers and sellers can submit shouts at any time in a specified trading round, and the market is cleared at a specific time and all possible transactions will be executed. Currently, there are broadly two typical types of double auctions, which are Continuous Double Auction (CDA, in which potential transactions are executed when a new shout arrives) and Clearing House (CH, in which transactions are executed until all traders have submitted their shouts). In Section 2.3.2.2, we will describe them in detail.

Intuitively, a double auction is a typical type of marketplace as we introduced in Section 2.2. Buyers and sellers in double auctions will have limit and cost prices respectively. However, in the double auction marketplace, the only known information is the bids submitted by buyers and asks submitted by sellers, which typically do not correspond to the actual limit and cost prices. This means that the equilibrium price and equilibrium quantity cannot be determined, and the optimal allocation cannot be obtained in double auctions (see Section 2.2).

Despite this fact, in theory, the competition of profit-motivated traders will eventually drive transaction prices to converge to the equilibrium price (which is consistent with the market theory detailed in Section 2.2). The reason is that, in order to remain competitive, the buyers (sellers) need to raise (lower) their prices when there is excess demand (supply), and thus cause transaction prices to move towards the equilibrium price. Such an hypothesis is confirmed by experimental analysis conducted by Smith (1962), where he stated:

> "The most striking general characteristic of [this] test ... is the remarkably strong tendency for exchange prices to approach the predicted equilibrium price for each of these markets. As the exchange process is repeated ... the variation in exchange prices tends to decline, and to cluster more closely around the equilibrium."

In his experiments, a group of (human) traders were split into two groups: buyers and sellers. Each trader could buy or sell one or more units of a homogenous commodity at a price, which is no lower than the given cost prices for sellers and no higher than the given limit prices for buyers. The transaction price was set at the average of the buyer's bid and the seller's ask.

In order to measure the convergence of transaction prices to the equilibrium price, Smith introduced a *coefficient of convergence*, α, given the history of H transaction prices TP_h ($h \in \{1...H\}$), which is expressed as a percentage and given by equation:

$$\alpha = \frac{\sqrt{\frac{1}{H}\sum_{h=1}^{H}(\mathrm{TP}_h - p^*)^2}}{p^*} \times 100 \tag{2.5}$$

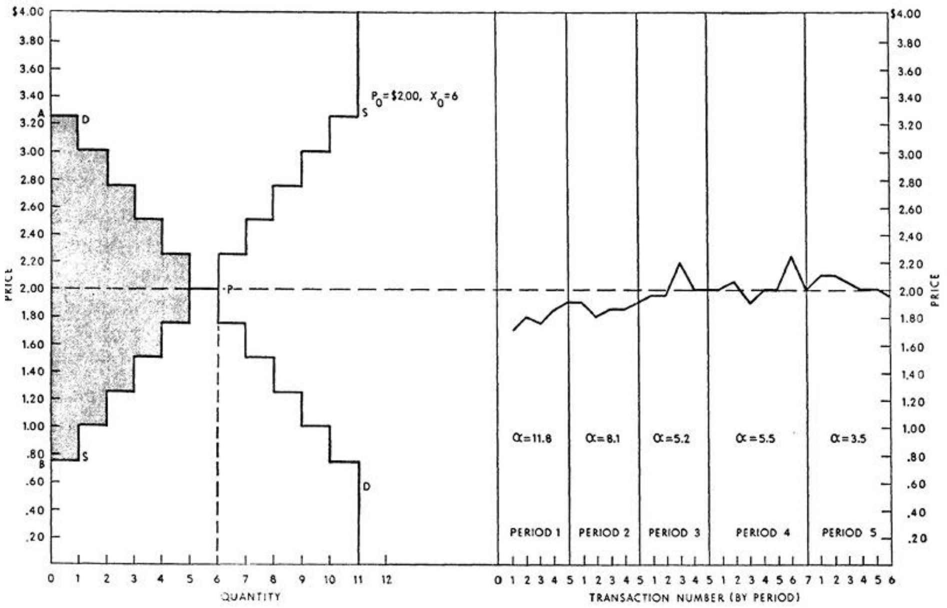

FIGURE 2.2: The left panel shows the underlying demand and supply of the marketplace. The dotted horizontal line indicates the equilibrium price, and the shaded region indicates the intra-marginal region of the marketplace, its area is equal to the theoretical overall profits. The right panel shows the history of transaction prices with α during the successive trading days (from Smith (1962)).

where TP_h is the price of the h-th transaction and p^* is the equilibrium price. When α decreases as traders make transactions, then we can see transaction prices indeed converge to the equilibrium price.

From Smith's experimental results (see Figure 2.2), we can see that transaction prices are close to the equilibrium price, and α decreases over time as the prices converge to the equilibrium price. Thus, Smith drew the conclusion that, when there is no collusion and all traders' shouts and transaction prices are public, the convergence of transaction prices to the equilibrium price can be reached in the marketplace. Smith's work also suggests that double auction marketplaces are bound to be efficient irrespective of the way that traders bid, and can achieve close to optimal allocative efficiency.

So far we have introduced the basics of double auctions. Research work about double auctions also includes bidding strategies of traders and market policies used by double auctions. As we will see, some existing bidding strategies are adopted in the specific context of the CAT competition. Since we will evaluate our competing marketplace in this context, it is necessary to introduce these bidding strategies. Furthermore, existing market policies will be the foundations of designing market policies for competing marketplaces. Therefore, we also need to introduce existing work about market polices. In addition to market policies, as we introduced in Section 1.1, the marketplace also needs a charging strategy to determine its fees. In the following, we will describe existing work in these three areas in turn.

2.3.2.1 Bidding Strategies and Equilibrium Analysis

In double auction marketplaces, traders need to use bidding strategies to determine the shouts submitted in the marketplaces. Existing literature describe a number of bidding strategies for double auctions—including the truth-telling bidding strategy in which traders submit their limit prices (resp. cost prices) as bids (resp. asks), the Fuzzy Logic based bidding strategy in which traders use heuristic fuzzy rules and fuzzy reasoning mechanisms in order to determine the best bid or ask given the current state of the marketplace (He et al., 2003), the adaptive-aggressiveness bidding strategy in which traders adapt their bids or asks based on a short-term and a long-term learning (Vytelingum et al., 2008), and the Q-strategy in which traders uses a Q-learning reinforcement learning approach to adapt their bids (Borissov, 2009; Borissov et al., 2010). However, in this thesis, since we choose to evaluate our design of a competing marketplace in the context of the CAT competition and these bidding strategies are not used in this context, we will not discuss them here. Specifically, in the following, we will introduce the bidding strategies used in the CAT competition, which are ZI-C, ZIP, GD and RE. Note that all these bidding strategies are heuristic based. Thus they can be widely used in different auction formats with different settings.

ZI-C Strategy:

Gode and Sunder (1993) developed a simple yet powerful strategy, called the *Zero-Intelligence* (ZI) bidding strategy. Here, zero intelligence means the trader is not motivated to pursue profit. Specifically, when the trader wants to submit a shout, it just selects a bid or an ask from a uniform distribution over a given range. In their experiments on allocative efficiency and equilibrium formation, Gode and Sunder considered two kinds of ZI strategies: the constrained ZI strategy (ZI-C) and the unconstrained ZI strategy (ZI-U). The former is restricted by budget constraints, which means that traders cannot trade at loss, whereas the latter is allowed to make loss-making transactions. Specifically, a ZI-C buyer draws a bid from a uniform distribution between the minimum allowed bid and its limit price, and the ZI-C seller draws an ask from a uniform distribution between its cost price and the maximal allowed ask. For ZI-U traders, their shouts are drawn from a uniform distribution between the minimum and the maximum allowed price in the marketplace.

The results of the simulations of marketplaces with ZI-C traders and ZI-U traders are shown in Figure 2.3. From these, we can see the differences between the resulting transaction prices when using ZI-C, ZI-U and human traders. In particular, we can see that the marketplace with ZI-U traders showed no evidence that transaction prices converge to an equilibrium price. In human marketplaces, on the other hand, transaction prices converge to the equilibrium price, and this is consistent with the classical microeconomic theory (see Section 2.2). However, they also found that, in the marketplace with ZI-C traders, there is a slow convergence to the equilibrium price during each trading day. They explained this convergence in the following manner. They assume that the buyers with the highest limit prices and the sellers with lowest cost prices have a greater chance to trade first. Then, the demand and supply curve shifts to the left as each good is traded, until only extra-marginal traders remain. As the demand and supply curve shifts, the range of feasible transaction prices narrows and transaction prices converge to the equilibrium price.

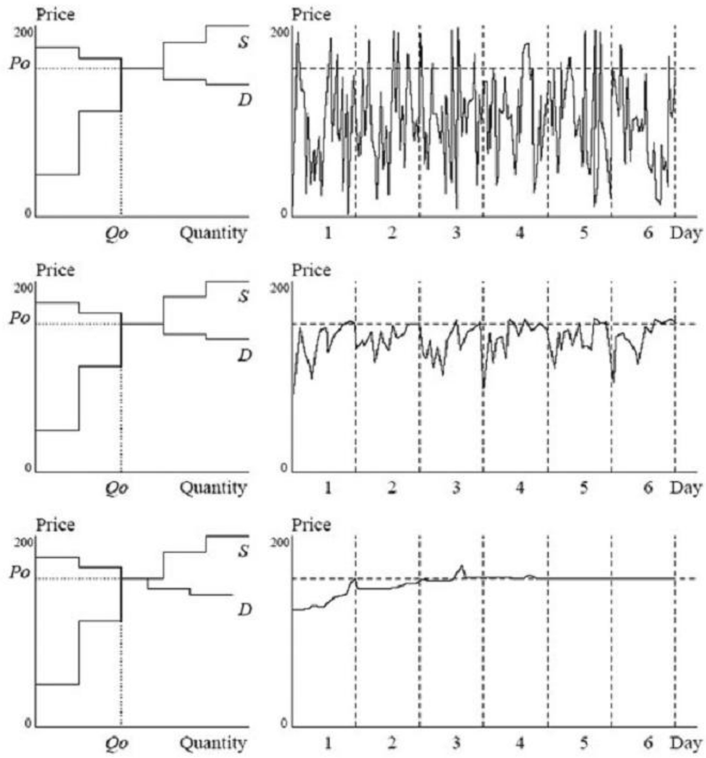

FIGURE 2.3: Result from one of Gode and Sunder's experiments. Figures on the left show the demand and supply used in the experiments. Figures on the right show transaction prices in different marketplaces where the top is the marketplace with ZI-U traders, the middle is with ZI-C traders and the bottom is with human traders (from Gode and Sunder (1993)).

The simulations also show that the allocative efficiency can be very high in the double auction marketplace even though traders are not profit-motivated and have no intelligence. Furthermore, they found that the efficiency of marketplaces with ZI-C traders is close to the efficiency of marketplaces with human traders. Thus they concluded that allocative efficiency is determined by the market structure (i.e. market policies), not by the bidding strategies of traders. However, they also pointed out that individual performance (i.e. an individual trader's profit) might be sensitive to individual intelligence.

In this paper, Gode and Sunder made a significant contribution to show that allocative efficiency is determined by market structure, and individual profit is determined by individual intelligence. However, their conclusion that ZI-C traders' transaction prices converge to the equilibrium price is attacked by Cliff and Bruten (1997), who showed that if demand and supply are not symmetric, the average transaction prices can be significantly different from the equilibrium price. To address this problem, they introduced the *Zero-Intelligence Plus* strategy which we describe in the next subsection.

ZIP Strategy:

The *Zero-Intelligence Plus* (ZIP) strategy was designed by Cliff and Bruten (1997) to show that more than zero intelligence is required to achieve efficiency close to that of marketplaces with human traders. This also attacked Gode and Sunder's claim that ZI-C traders could achieve and stabilise at equilibrium in double auction marketplaces. First, we give a simple description of the ZIP strategy, and then describe the experimental results to show ZIP's performance.

In more detail, the ZIP strategy uses a history of market information, and adjusts the trader's profit margin $\mu(t)$ according to the future market conditions in order to remain competitive. Here, the profit margin determines the difference between the trader's limit price (cost price) and the bid (ask). According to the profit margin, the seller i submits its ask $a_i(t)$ at the time t:

$$a_i(t) = \lambda_i(1 + \mu_i(t)) \tag{2.6}$$

where λ_i is the cost price of seller i. Similarly, a ZIP buyer j's bid at the time t is:

$$b_j(t) = \lambda_j(1 - \mu_j(t)) \tag{2.7}$$

where λ_j is the limit price of the buyer. The ZIP sellers raise or lower their profits by increasing or decreasing $\mu_i(t)$, and similar for buyers. By dynamically modifying $\mu_i(t)$ or $\mu_j(t)$, sellers or buyers remain competitive against other traders in the marketplace. The adaptation of the profit margin is based on the simple Widrow-Hoff learning algorithm (Widrow and Hoff, 1960). The actual adaption rules are shown in Figure 2.4.

Adaptive Rules for the ZIP Seller:

if (last shout was accepted at price $q(t)$)
  1. any seller i for which $a_i(t) \leq q(t)$ should raise its profit margin
  2. if (last shout was a bid)
    1. any active seller i for which $a_i(t) \geq q(t)$ should lower its margin
else
  2. if (last shout was an ask)
    1. any active seller i for which $a_i(t) \geq q(t)$ should lower its margin

Adaptive Rules for the ZIP Buyer:

if (last shout was accepted at price $q(t)$)
  1. any buyer j for which $b_j(t) \geq q(t)$ should raise its profit margin
  2. if (last shout was an ask)
    1. any active buyer j for which $b_j(t) \leq q(t)$ should lower its margin
else
  2. if (last shout was a bid)
    1. any active buyer j for which $b_j(t) \leq q(t)$ should lower its margin

FIGURE 2.4: The ZIP trading strategy.

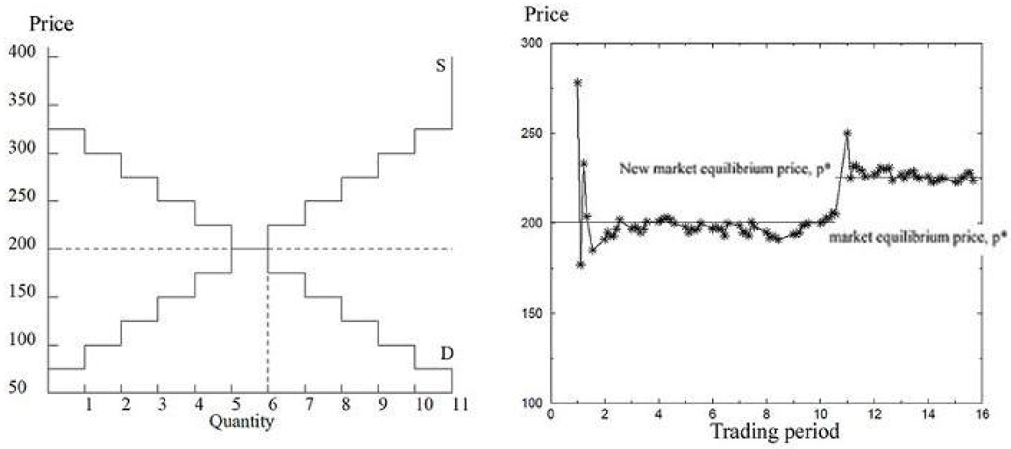

FIGURE 2.5: The left panel illustrates the demand and supply used for the first 11 trading periods, and then the equilibrium price is increased to 225. The right panel shows the results from simulations with ZIP traders, $*$ means price of each transaction (from Cliff and Bruten (1997)).

The simulation results using ZIP strategies are shown in Figure 2.5. We can see that the transaction prices converge towards the equilibrium prices after just a few days, and remain at that level with low variance. Cliff also showed that the profit dispersion of ZIP traders was significantly lower than that of ZI-C traders. Furthermore, by considering a sudden change in endowment of limit and cost prices to buyers and sellers respectively at the beginning of period 12, where the demand and supply changed and the equilibrium price increased from 200 to 225, they showed that the transaction prices can rapidly converge to the new equilibrium price. This means this strategy has a capacity to respond quickly to the changing market conditions.

GD Strategy:

Another well-known intelligent bidding strategy is called GD (Gjerstad and Dickhaut, 1998). This strategy is based on a belief function that indicates the probability of a particular bid(ask) being accepted in the marketplace. In more detail, the traders form their beliefs according to the history of marketplace data, and particularly on the frequencies of submitted bids and asks and of accepted bids and asks resulting in transactions. Based on the belief function, the GD trader submits a shout which can maximise its expected profit, i.e. the product of its belief function and its profit if a transaction occurs.

In the GD strategy, the seller's belief function, $\hat{p}(a)$, is constructed as follows: if an ask $a' < a$ has been rejected by the marketplace, then the ask a will also be rejected. Similarly, if an ask $a' > a$ has been accepted, than the ask a will also be accepted. Furthermore, if a bid $b' > a$ is accepted, then an ask $a' = b' > a$ would have been accepted, and then a will be accepted by the marketplace. Similarly, the buyer's belief function $\hat{a}(b)$ can be constructed. In what follows, we define the bid and ask frequencies $\forall d \in D$, where D is the set of all permissible shout prices in the marketplace.

Bid Frequencies: $\forall d \in D$, $B(d)$ is the total number of bids submitted at price d, $\mathrm{TB}(d)$ is the frequency of accepted bids at d, and $\mathrm{RB}(d)$ is the frequency of rejected bids at d.

Ask Frequencies: $\forall d \in D$, $A(d)$ is the total number of asks submitted at price d, $\mathrm{TA}(d)$ is the frequency of accepted asks at d, and $\mathrm{RA}(d)$ is the frequency of rejected asks at d.

Now, the seller's belief function for each possible ask price a is given by:

$$\hat{p}(a) = \frac{\sum_{d \geq a} \mathrm{TA}(d) + \sum_{d \geq a} B(d)}{\sum_{d \geq a} \mathrm{TA}(d) + \sum_{d \geq a} B(d) + \sum_{d \leq a} \mathrm{RA}(d)} \tag{2.8}$$

and the buyer's belief function for each possible bid price b is given by:

$$\hat{q}(b) = \frac{\sum_{d \leq b} \mathrm{TB}(d) + \sum_{d \leq b} A(d)}{\sum_{d \leq b} \mathrm{TB}(d) + \sum_{d \leq b} A(d) + \sum_{d \geq b} \mathrm{RB}(d)} \tag{2.9}$$

Moreover, the seller and buyer's belief function is modified to satisfy the NYSE shout accepting policy[5]. That is, the belief function on the ask, which is higher than the current outstanding ask O_{ask}, is set to 0 and cannot be accepted, and when the bid is lower than the outstanding bid O_{bid}, then belief function on this bid is set to 0, and cannot be accepted.

In addition, because the belief function is defined on the set of all bids and asks, then we need to extend the beliefs to the space of all potential bids or asks, which are constrained by the outstanding bid O_{bid} and outstanding ask O_{ask} and the step-size of the belief function. To this end, *Cubic spline interpolation* is used on each successive pair of data points to calculate the belief of points in between them. In particular, a cubic function, $p(a) = \alpha_3 a^3 + \alpha_2 a^2 + \alpha_1 a + \alpha_0$, is constructed with the following properties:

1. $p(a_k) = \hat{p}(a_k)$
2. $p(a_{k+1}) = \hat{p}(a_{k+1})$
3. $p'(a_k) = 0$
4. $p'(a_{k+1}) = 0$

where a_k and a_{k+1} are the successive ask prices. $p'(a_k)$ is the first derivative of $p(a_k)$. The coefficients, α_i, satisfying the above properties, are given by the solution to the following equation:

$$\begin{bmatrix} a_k^3 & a_k^2 & a_k & 1 \\ a_{k+1}^3 & a_{k+1}^2 & a_{k+1} & 1 \\ 3a_k^2 & 2a_k & 1 & 0 \\ 3a_{k+1}^2 & 2a_{k+1} & 1 & 0 \end{bmatrix} \begin{bmatrix} \alpha_3 \\ \alpha_2 \\ \alpha_1 \\ \alpha_0 \end{bmatrix} = \begin{bmatrix} \hat{p}(a_k) \\ \hat{p}(a_{k+1}) \\ 0 \\ 0 \end{bmatrix} \tag{2.10}$$

Now we have constructed the sellers' belief function, the buyer's belief function, $q(b)$, can be constructed similarly using the pairs $(b_k, \hat{q}(b_k))$ and $(b_{k+1}, \hat{q}(b_{k+1}))$. After defining the belief function, Gjerstad and Dickhaut proved (see Gjerstad and Dickhaut (1998) for the proof) that

[5]The NYSE shout accepting policy requires that a submitted bid is higher than the outstanding bid (i.e. the current highest unmatched bid in the marketplace), and a submitted ask is lower than the outstanding ask (i.e. the current lowest unmatched ask in the marketplace) respectively.

the beliefs are monotonically non-decreasing(non-increasing) for bids(asks). In particular, from Figure 2.6, we can see the belief that an ask $a > a'$ is accepted is lower than the belief of a' being accepted, and, similarly, the belief of a bid that $b < b'$ is accepted is lower than that of b.

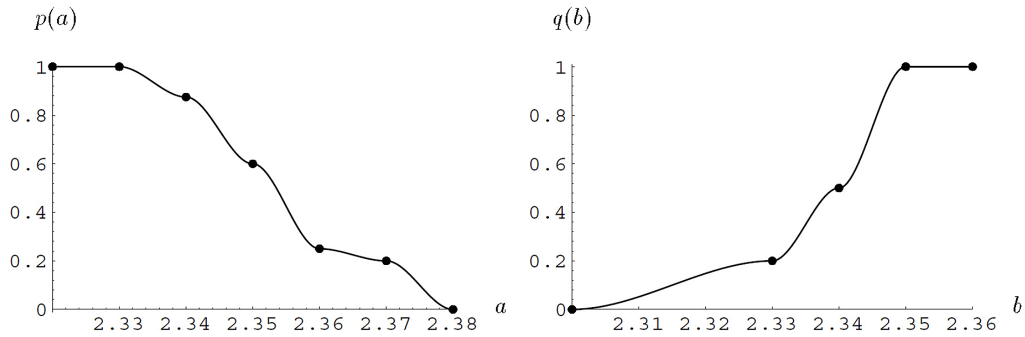

FIGURE 2.6: A typical belief function of a buyer (right) and seller (left) using GD strategy (from Gjerstad and Dickhaut (1998)).

After getting the belief function, the GD trader can form its bid or ask that maximises its expected profit. The expected profit is defined as the product of the trader's belief function and its utility function. The utility function equals the difference between the seller i's ask and its cost price c_i, or the difference between the buyer j's bid and its limit price l_j (i.e. the trader's hidden profit through shading). We can see that as a trader increases its bid or decreases its ask, its utility increases. However, from Figure 2.6, we can see that its belief of shout acceptance decreases. Thus the trader has to make a trade-off between increasing its belief of shout acceptance and increasing its utility. In more detail, the trader's utility function and its profit maximisation is given as follows:

For a seller i,

$$U(a) = \begin{cases} a - c_i & \text{if } a > c_i \\ 0 & \text{if } a \leq c_i \end{cases}$$

For a buyer j,

$$U(b) = \begin{cases} l_j - b & \text{if } b < l_j \\ 0 & \text{if } b \geq l_j \end{cases}$$

$$a^* = \arg \max_{a \in (o_{ask}, o_{bid})} [U(a).\hat{p}(a)] \tag{2.11}$$

$$b^* = \arg \max_{b \in (o_{ask}, o_{bid})} [U(b).\hat{q}(b)] \tag{2.12}$$

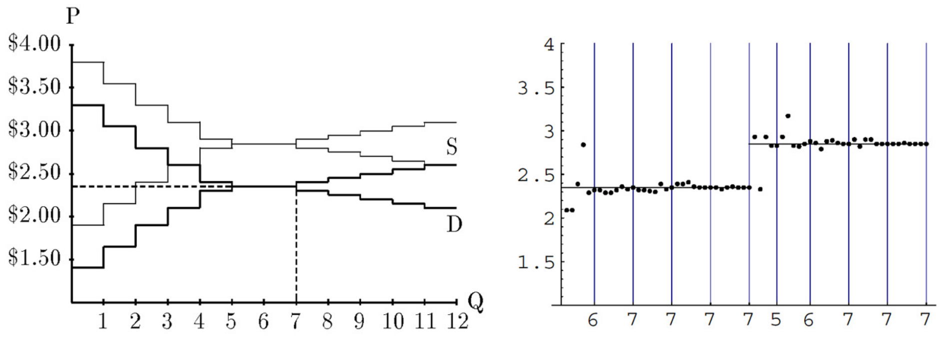

FIGURE 2.7: Left panel illustrates demand and supply of the marketplace. Note the change in demand and supply after 5 trading days. Results of the simulations with GD traders are shown in the right panel. The x-axis is divided into the different trading days, with the x-axis values corresponding to the number of transactions on each trading day. The y-axis values correspond to the transaction prices (from Gjerstad and Dickhaut (1998)).

Again they carried out a series of simulations to evaluate the performance of this strategy. In particular, the results of the simulations showed that the allocative efficiency of marketplaces with GD traders was close to optimal, and transaction prices converged rapidly to the equilibrium price (see Figure 2.7). By shifting the demand and supply after several trading days, it was also shown that GD traders responded quickly to the changing market conditions, and transaction prices quickly converged to the new equilibrium price.

RE Strategy:

We now introduce another bidding strategy called Roth-Erev (RE), which attempts to mimic human game-playing behaviour. In a series of studies, Roth and Erev attempted to understand how people learn individually to behave in games with multiple strategic players. To this end, they developed a myopic reinforcement learning algorithm, referred to as the RE algorithm (Roth and Erev, 1998; Nicolaisen et al., 2001). We now describe this strategy in more detail. The authors suppose that there are G feasible actions (i.e. bids or asks) for each trader in the marketplace. At the initial trading round, each trader i assigns an equal *propensity* $\phi_{ig}(1)$ to each feasible action g given by $\phi_{ig}(1) = \frac{\gamma(1)X}{G}$, where X is the average profit that traders can achieve in any given trading round, and $\gamma(1)$ is a scaling parameter. Moreover, each trader i assigns an initial equal choice probability $\psi_{ig}(1)$ to each of its feasible actions g, given by $\psi_{ig}(1) = \frac{1}{G}$. Here, "1" means the initial trading round. The trader i then probabilistically selects a feasible action g' to submit according to its current choice probability. If the shout g' can be matched, then a transaction is executed and the trader i gets a profit $R(i, g', 1)$. Now we suppose that the trader i is at the end of the nth trading round and in this round, the trader has submitted a feasible action g' and achieved a profit $R(i, g', n)$. The trader i then updates its existing action propensities $\phi_{ig}(n)$ based on its newly earned profit as follows. Given any feasible action g, the propensity $\phi_{ig}(n+1)$ for choosing g in the next trading round $n+1$ is determined as:

$$\phi_{ig}(n+1) = (1-r)\phi_{ig}(n) + \rho(i, g, g', n, G, e) \tag{2.13}$$

where r is the recency parameter, e is the experimentation parameter and $\rho(\cdot)$ is the update function reflecting the experience gained from past trades. The recency parameter r slowly reduces the importance of past experience. The update function $\rho(\cdot)$ is given by

$$\rho(i, g, g', n, G, e) = \begin{cases} R(i, g', n)(1 - e) & \text{if } g = g' \\ \frac{R(i,g',n)e}{(G-1)} & \text{if } g \neq g' \end{cases}$$

The selected action g' is reinforced or discouraged based on the profit $R(i, g', n)$ earned subsequent to this selection. Given the updated propensities $\phi_{ig}(n + 1)$ for trading round $n + 1$, the trader i updates choice probability $\psi_{ig}(n + 1)$ for its feasible actions g in the trading round $n + 1$ given by the equation:

$$\psi_{ig}(n + 1) = \frac{\phi_{ig}(n + 1)}{\sum_{g=1}^{g=G}(\phi_{ig}(n + 1))} \tag{2.14}$$

Finally the trader i selects a feasible action according to its current choice probability.

To date, there exists no literature on analysing the convergence of transaction prices to the equilibrium price when traders adopt the RE strategy. However, in Section 5.1, we analyse this based on the JCAT platform, and find that transaction prices of traders adopting RE strategy do not converge to the equilibrium price.

Equilibrium Analysis:

In the above, we have introduced four bidding strategies, and these strategies were usually evaluated in the homogeneous environment where all traders use the same strategy. However, in double auctions, traders may be able to use different bidding strategies, which results in the problem of determining which bidding strategy traders should choose and what is the Nash equilibrium bidding strategy. In this context, Phelps et al. (2006) used evolutionary game theory (EGT) to analyse the equilibrium bidding strategies when traders can use three different bidding strategies: truth-telling (TT), RE, and PvT (a ZIP-like strategy, which is modified for persistent-shout marketplaces (Preist and van Tol., 2003)) in two different types of double auction marketplaces (Clearing House and Continuous Double Auction, in the following section, we will introduce these two double auctions in detail). They show that in the Clearing House with 6 traders, in equilibrium, traders will have 50% probability of using RE strategy and 50% probability of using PvT strategy, and in the case with 8 traders, traders have 17% probability of using TT strategy, 18% probability of using RE strategy and 65% probability of using PvT strategy. Furthermore, in the Continuous Double Auction, in equilibrium, traders will have 100% probability of using RE strategy in both cases with 6 traders and 8 traders. Moreover, in Phelps et al. (2010), they also extend their work to the case with four different bidding strategies: TT, RE, GD and TK (Kaplan's sniping strategy (Friedman and Rust, 1993)). Moreover, Vytelingum et al. (2008) used EGT to evaluate their AA strategy against ZIP or GDX (which is a improved bidding strategy based on GD (Tesauro and Bredin, 2002)), and showed that in equilibrium, traders are more likely to use the AA strategy. Note that although the above works analyse the Nash equilibrium of existing bidding strategies, they fail to answer what exactly traders will bid in equilibrium. Furthermore, they do not analyse how market fees affect traders' bidding behaviour. In our work, we will use fictitious play to derive what shouts traders will submit in equilibrium and how different types of market fees affect their equilibrium bidding strategies (see Section 4.3.1).

2.3.2.2 Market Policies

After describing the key bidding strategies, we now review related work about market policies for double auctions. As we introduced in Section 1.1, in the design of a double auction, we need to specify four market policies, which are timing policy (determining when to clear the market), matching policy (determining how to match buyers with sellers), pricing policy (determining the transaction price of the matched buyer and seller) and shout accepting policy (determining whether to admit the shout placed by the trader or not). In the following, we will describe the related work about these policies respectively.

Timing Policy:

A timing policy determines when to clear the market (i.e. when to match bids with asks to make transactions). There are a number of alternatives. The first is to collect all bids and asks and to clear the market at the end of the trading day in order to maximise profits. A double auction adopting such an approach is often referred to as a Clearing House (CH). Another approach is to continuously clear whenever a bid or ask is accepted in the marketplace. This kind of double auction is often called a Continuous Double Auction (CDA). Phelps et al. (2006) used evolutionary game theory to compare the efficiencies of these alternatives. In their experiments, traders are free to choose the following bidding strategies: truth-telling (TT), RE, and PvT. The simulation results are shown in Tables 2.1 and 2.2, where the sum of traders' profits is equal to the allocative efficiency of the marketplace since the sum of expected profits is 1. The insights obtained are as follows. In the CH, we can see that the most likely strategy to be used by traders is the PvT strategy, and in the CDA, the RE strategy is the most likely one to be used. We also see that the total profits of all traders under the truth-telling strategy in the CDA are relatively low, only 0.86 in this case (mean of 0.87 in Table 2.1 and 0.85 in Table 2.2). This might suggest that the CDA has low allocative efficiency, only 86%. However, because the RE strategy outperforms all other strategies, all traders will choose it eventually. Then, in this case, the profit is 0.98. In the CH, we can see that each strategy results in the same profit, 1, and can conclude that the CH will yield 100% allocative efficiency in all cases. Now, although the CDA may yield lower efficiency, it is still widely used in real exchange marketplaces since it is better able to handle larger volumes of trades by clearing the market quickly (Friedman and Rust, 1993). Moreover, switching to a CDA from a CH, as the NYSE did in the 1860s, does not seem to cause a large loss of efficiency in practice (actually, from the experiment we can see when all traders adopt the RE strategy, the CDA is quite efficient with 98% allocative efficiency).

In sum, the CDA is better for higher numbers of transactions since it clears the market whenever

Equilibrium	CH probability	profit(allocative efficiency)	CDA probability	profit(allocative efficiency)
TT	0.00	1.00	0.00	0.87
RE	0.50	1.00	1.00	0.98
PvT	0.50	1.00	0.00	0.93

TABLE 2.1: Equilibrium and profit for 6 traders.

Equilibrium	CH probability	profit(allocative efficiency)	CDA probability	profit(allocative efficiency)
TT	0.17	1.00	0.00	0.85
RE	0.18	1.00	1.00	0.98
PvT	0.65	1.00	0.00	0.92

TABLE 2.2: Equilibrium and profit for 8 traders.

a new bid or ask arrives. This means traders can buy or sell items quickly, but the CDA is worse for allocative efficiency since some extra-marginal traders may steal a deal. CH is better for allocative efficiency since it clears the market when all traders have submitted shouts. This means intra-marginal buyers will most likely be matched with intra-marginal sellers, but it is worse in terms of the number of transactions since it cannot clear the market quickly in practice. Given this, it is clear that, when determining when to clear the market, we need to consider the trade-off between maximising profits and maximising the number of transactions.

Matching Policy:

A matching policy determines how to clear the market (i.e. how to match bids with asks to make transactions). In this context, *equilibrium matching* (ME) is the most commonly used one (Niu et al., 2008a). This policy will match intra-marginal buyers with intra-marginal sellers, and the transaction prices are set as the equilibrium price. Since intra-marginal traders can make transactions to guarantee their profits, ME achieves a high allocative efficiency. Note that when using this matching policy, since the transaction prices are set at the equilibrium price, the marketplace can arbitrarily match successful bids and asks to each other. However, in our theoretical analysis (see Chapters 3 and 4), since we use the k-pricing policy with $k = 0.5$ to set the transaction prices (see the below), we need to operate the equilibrium matching in a specific way. In more detail, we will match buyers having the v-th highest bids with sellers having v-th lowest asks. By so doing, we still guarantee intra-marginal traders' profits. Another matching policy, named *Max-volume matching* (MV) (Niu et al., 2008a), aims to increase transaction volume based on the observation that a high intra-marginal bid can be matched with an extra-marginal ask which is a little lower than the high bid. The MV matching policy can increase the number of transactions, but at the same time can cause a loss of profits, and thus result in a low allocative efficiency. Under this situation, a trade-off is also needed between maximising profits (using ME) and maximising the number of transactions (using MV).

Pricing Policy:

The pricing policy determines the transaction prices for matched bids and asks. This means that the pricing policy redistributes the profits among the traders and thus influences the profits of the buyers and sellers. When designing the pricing policy, we need to consider the allocative efficiency, buyers' efficiency and sellers' efficiency[6]. Furthermore, the degree of transaction price fluctuation also needs to be considered (Niu et al., 2006). This is because with reduced price fluctuations, transaction prices are guaranteed to be close to the equilibrium price. With this guarantee for a fair transaction price, more traders are likely to prefer to join this marketplace, which should, in turn, make the marketplace more profitable.

To this end, Phelps et al. (2003) consider how to design a pricing rule for the CDA to maximise V, which is the combination of allocative efficiency, the buyers' efficiency and the sellers' efficiency:

$$V = \frac{\text{allocative efficiency}}{2} + \frac{\text{buyers}' \text{ efficiency}}{4} + \frac{\text{sellers}' \text{ efficiency}}{4} \tag{2.15}$$

They propose two approaches to designing the pricing policy. First, the transaction price is calculated according to:

$$\text{TP} = k \cdot p_a + (1 - k) \cdot p_b \quad (k \in [0, 1]) \tag{2.16}$$

where p_a is the matched ask and p_b is the matched bid. This is called the k-pricing policy and is a traditional pricing policy used in double auction marketplaces. In order to maximise V, they want to optimise k. To this end, their simulation results are shown in Figure 2.8, and as we can see the best value for k in this setting is 0.5.

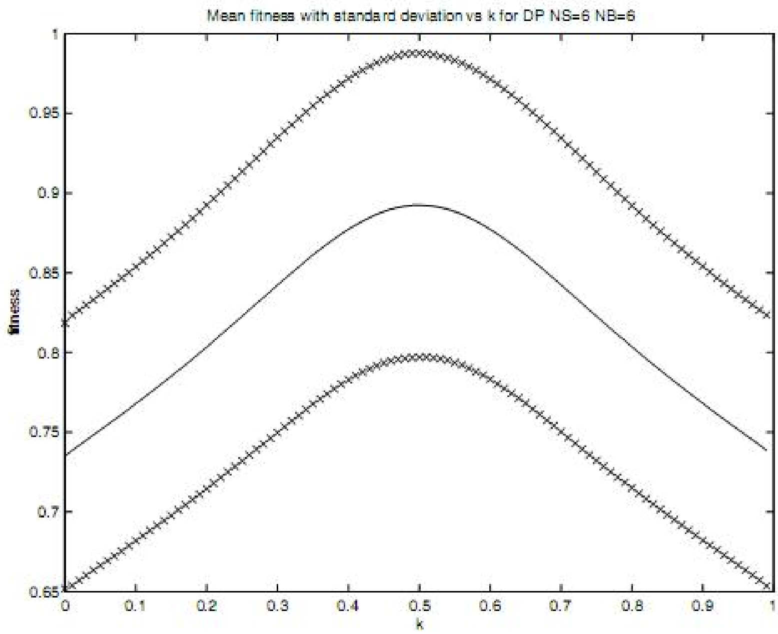

FIGURE 2.8: Fitness V (with standard deviation) plotted against k for a marketplace with 12 traders (from Phelps et al. (2003)).

[6] This is the ratio of actual profits of buyers(sellers) to their expected profits when transactions are executed at the equilibrium price.

The second approach considers all possible combinations of the buyer's bid p_b and the seller's ask p_a. The pricing policy was then allowed to evolve using genetic programming (see Phelps et al. (2003) for details), and finally it was approximately equal to $0.5p_b + 0.5p_a$, apart from a small variation when the ask is small or when the difference between the bid and ask is marginal.

As can be seen, both approaches show that the k-pricing policy ($k = 0.5$) is efficient both in terms of allocative efficiency and the traders' efficiency.

In terms of reducing price fluctuations, Niu et al. (2006) explore modifying the traditional policy, referred to as the k-pricing (given by Equation 2.16) to the n-pricing policy. The n-pricing policy keeps a sliding window of size n of matching pairs of bids and asks which are used to set the transaction prices. It is given by:

$$\mathrm{TP} = \frac{1}{2n} \sum_{z=T-n+1}^{T} (p_{a_z} + p_{b_z}) \quad (p_{aT} \leq \mathrm{TP} \leq p_{bT}) \tag{2.17}$$

where p_{b_z} and p_{a_z} are the accepted bid and ask corresponding to the z-th transaction respectively, T is the latest transaction and TP is the price at which the transaction is set using the n-pricing policy. Here, TP is bounded between the buyer's bid p_{bT} and the seller's ask p_{aT}. Note that the n-pricing policy becomes the k-pricing policy with $k = 0.5$ when $n = 1$, and the auction is then a traditional CDA. On the other hand, when n is equal to the total number of transactions, the n-pricing policy is the rule commonly used in a CH. Thus, n ranges over a continuous space of double auctions, with the Continuous Double Auction and the Clearing-House Double Auction at either end.

The authors then compare the price fluctuations of n-pricing policy with $n = 4$ with k-pricing policy with $k = 0.5$ in the CDA by measuring the coefficient of convergence α given a simple and non-intelligent behaviour (with ZI-C traders) and a more complex and intelligent behaviour (with GD traders). As an example, the marketplace named kCDA-ZIC means the CDA with k-pricing policy and populated by ZI-C traders. The results, given in Figure 2.9 and Table 2.3, showed that the price fluctuations were reduced in both marketplaces which use n-pricing policy, with the better improvement in the marketplace with ZI-C traders (indicated by the relatively larger decrease in α). This is because ZI-C traders randomly submit bids and asks, and the spread of matching bids and asks is typically quite high. On the other hand, the more intelligent GD traders submit bids and asks that tend to converge towards the equilibrium price, such that the spread is then greatly reduced, and the n-pricing policy is then only marginally more effective.

To sum up, the effect of reducing price fluctuations by using the n-pricing policy is restricted by traders' bidding strategies. Moreover, there is no evidence to show that the allocative efficiency improves when a marketplace adopts the n-pricing policy, and in some cases, the allocative efficiency even decreases.

Shout Accepting Policy:

The shout accepting policy considers which bids and asks should be accepted by the market-

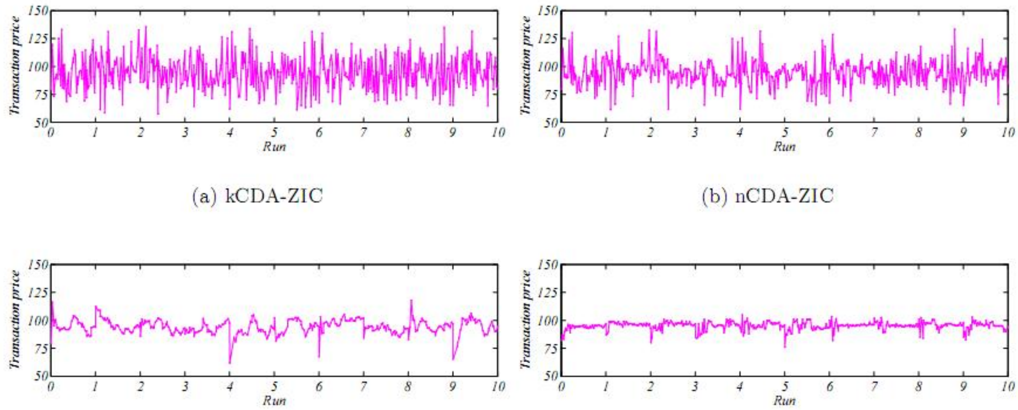

(a) kCDA-ZIC (b) nCDA-ZIC

(c) nCDAEE-ZIC (d) kCDA-GD

FIGURE 2.9: Transaction prices, plotted as 10 different runs for the k-pricing policy and for the n-pricing policy (from Niu et al. (2006)).

Auction Type	α mean	α stdev	E_a mean	E_a stdev	Auction Type	α mean	α stdev	E_a mean	E_a stdev
kCDA-ZIC	15.613	1.352	95.732	1.537	kCDA-GD	4.198	1.140	99.224	0.548
nCDA-ZIC	**12.180**	1.532	**95.732**	1.537	nCDA-GD	**4.000**	1.1221	**99.224**	0.548
kCDAEE-ZIC	**9.360**	1.039	93.296	2.061	kCDAEE-GD	5.374	4.909	84.612	26.165
nCDAEE-ZIC	***6.028***	1.353	93.788	1.938	nCDAEE-GD	5.271	5.245	83.460	26.783
nCDAEEd5-ZIC	**6.221**	1.300	**96.072**	1.657	nCDAEEd5-GD	4.242	2.779	95.996	13.126
nCDAEEd10-ZIC	**6.879**	1.117	***96.868***	1.187	nCDAEEd10-GD	**4.054**	1.262	98.768	1.092
nCDAEEd15-ZIC	**8.000**	1.017	**96.864**	1.179	nCDAEEd15-GD	**4.016**	1.256	99.132	0.684
nCDAEEd20-ZIC	**9.220**	1.002	**96.744**	1.144	nCDAEEd20-GD	***3.994***	1.217	***99.224***	0.548

TABLE 2.3: Metrics for KCDAs, nCDAs and nCDAEEs measured over 10 trading days. Bold face indicates the corresponding marketplace outperforms or equals its traditional kCDA counterpart. ***Bold italic*** points out the best result in the corresponding ZI-C or GD marketplace group (from Niu et al. (2006)).

place. This policy is typically used to speed up the transaction process, and to reduce the price fluctuation. For example, when the marketplace only accepts bids considerably above the equilibrium prices and asks considerably below the equilibrium price, bids and asks are much more easily matched, and thus the transaction process is speeded up.

In more detail, the most common shout accepting policy is the NYSE shout accepting policy, which requires that a submitted bid or ask improves on the outstanding bid or the outstanding ask respectively. Such a policy is also called a *quote-beating accepting policy* (Niu et al., 2008a). In order to reduce price fluctuation, Niu et al. (2006) replaced the traditional NYSE shout accepting policy with a novel estimated equilibrium shout accepting policy (EE). Specifically, they estimated the equilibrium price, denoted by $\hat{p}^*$, using a sliding window of the n' latest transactions:

$$\hat{p}^* = \frac{1}{n'} \sum_{x=T-n'+1}^{T} \mathrm{TP}_x \qquad (2.18)$$

where TP_x is the transaction price calculated in equation 2.17. According to the estimated equilibrium price, the buyers have to bid above $(\hat{p}^* - \epsilon)$, and the sellers have to ask below

$(\hat{p}^* + \epsilon)$ $(\epsilon \geq 0)$. Since transaction prices deviate considerably from the equilibrium price at the beginning of trading, a parameter ϵ is introduced to relax the range of allowed bids and asks. The efficiency and price fluctuations with the new shout accepting policy (denoted by CDAEEϵ) are given in Table 2.3, where when $\epsilon = 0$, marketplaces are called kCDAEE-ZIC, nCDAEE-ZIC, kCDAEE-GD and nCDAEE-GD, and when $\epsilon = 5, 10, 15, 20$, marketplaces are called nCDAEEd5-ZIC, nCDAEEd5-GD, nCDAEEd10-ZIC, nCDAEEd10-GD, nCDAEEd15-ZIC, nCDAEEd15-GD, nCDAEEd20-ZIC, nCDAEEd20-GD. As an example, the marketplace named nCDAEEd15-GD means the CDA with n-pricing policy and estimated equilibrium shout accepting policy with $\epsilon = 15$ and populated by GD traders. Through simulations, a number of observations can be made. Firstly, they consider the case where $\epsilon = 0$. We can see α is then considerably smaller for nCDAEE-ZIC, although the new shout accepting policy decreases efficiency in that case. However, with GD traders, the performance is even worse. The price fluctuations actually increase and the allocative efficiency decreases. The authors conjectured that the GD trader adapts its bids or asks even though they are on the wrong side of the estimate and thus get rejected, and it has insufficient time to adapt sufficiently to be efficient. Because of this, the authors consider relaxing the shout accepting policy and consider different values for the parameter ϵ. The results showed that, the marketplaces with $\epsilon \geq 5$ were indeed more efficient with either ZI-C or GD traders, than with $\epsilon = 0$, although when they increased ϵ, the price fluctuations decreased with GD traders, but increased with ZI-C traders.

In summary, when the marketplace adopts the estimated equilibrium shout accepting policy, in some cases, its allocative efficiency indeed increases. Thus when we compare the different market policies, we will also consider the estimated equilibrium shout accepting policy.

2.3.2.3 Charging Strategy

In addition to establishing market polices, the double auction marketplaces also need to determine fees charged to participating traders. There are many kinds of fees charged to traders, such as a registration fee charged when traders enter the marketplace, or a transaction fee charged when a transaction is executed. By charging fees, marketplaces earn profits. In isolated marketplaces, charging strategies are typically not considered since traders have no other marketplace they can go to and thus fees have little impact on the behaviour of the traders. Currently, there is no literature particularly considering a charging strategy in an isolated double auction marketplace. However, works on charging strategies for competing marketplaces exist (see Section 2.4.2.1 and 2.5.1). Furthermore, even in an isolated marketplace environment, when traders can can choose whether or not to trade, how to set fees to maximise the market profit is also interesting. Specifically, the total profit extracted from traders depends on the number of traders participating in the marketplace and the fee charged to each trader. A higher fee causes more traders not to choose the marketplace (even though they have no other marketplaces to enter). Therefore, an appropriate fee should be set to maximise the profit. In 4.4.1, we will address this problem in an isolated marketplace environment.

2.4 Competing Marketplaces

In the previous section, we have introduced related work about isolated single-sided and double-sided auctions, and discussed both bidding strategies and market policies for these settings. However, as the development of global economy continues, more and more marketplaces emerge, and traders have greater choices of which marketplaces to select and trade in. To this end, in this section, we present existing work related to the competing marketplaces. In particular, when trading in multiple marketplaces, traders need to consider where to trade and how to trade, i.e. the market selection strategy and the bidding strategy. In the following, we will describe the related work about traders' strategies, in terms of both how they should bid (the bidding strategy) and which marketplace they should choose (the market selection strategy), and then we introduce the related work about competition between marketplaces.

2.4.1 Single-Sided Marketplaces

In this section, we describe the traders' strategies and marketplaces' strategies in the context of multiple single-sided auctions (i.e. English auctions, Dutch Auctions, etc), since as we mentioned before, research work about competing single-sided marketplaces will provide us some insights for analysing competing double auction marketplaces.

2.4.1.1 Traders' Strategies

In the environment with multiple competing single-sided auctions, in order to make effective trading decisions, traders need to determine the best set of auctions in which to bid, and determine how much to bid in the chosen auctions. Specifically, they often have to make decisions in a dynamic, unpredictable and time-constrained environments. Thus it is often difficult for traders to find an optimal strategy that can be used in practical contexts. For this reason, researchers mainly adopt empirical approaches to design trading strategies across multiple marketplaces.

To this end, Preist designed an algorithm for traders that purchase one or more identical goods from multiple English auctions which may terminate simultaneously or at different times (Preist et al., 2001b,a). When all auctions terminate simultaneously, the trader uses a coordination mechanism to ensure that he has the lowest leading bids to purchase the appropriate number of goods. Furthermore, when auctions terminate at different times, the trader first calculates its expected utility in the auction which is about to terminate, and compares it to expected utilities in the remaining non-terminating auctions (the calculation is based on the trader's beliefs about other traders' private values). Given the comparison, the trader must decide whether to place a higher bid in the auction which is about to terminate, or withdraw from it. How the traders update their beliefs of other traders' private values is in a similar spirit to the GD strategy, except that GD strategy is only applied in a single double auction marketplace. They also demonstrated that, when more traders adopt this algorithm, the allocative efficiency is increased.

In addition to Preist's work, Byde also proposed a dynamic programming approach for traders that participate in multiple English auctions to buy a single good (Byde, 2001a,b), and shows that the dynamic programming approach is more effective at obtaining high utilities compared to other algorithms. Based on this initial work, Byde et al. (2002) further developed a framework in which traders can make rational bidding decisions to buy multiple identical goods across multiple heterogenous single-sided auctions (English, Dutch, sealed-bid first-price and Vickrey auctions) with varying start and end times (i.e. the times for starting auctions and terminating auctions can be different). Such a framework aims to find the *optimal* decision making behaviour for traders. Furthermore, since finding the optimal decision making behaviour is infeasible given the time-constrained nature of auctions, in order to get a practical solution, Byde et al. (2002) also implemented a heuristic algorithm to approximate the decision making behaviour.

Moreover, He et al. (2006) developed a heuristic algorithm for traders to buy multiple identical goods from multiple English auctions with varying start and end times. This algorithm works by using a fuzzy neural network to predict auctions' closing prices and to decide the allocation of the goods to the customers, and then calculates the satisfaction degree of the allocation. Based on this, it calculates the set of auctions which it believes are the best to bid in. Moreover, rather than just bidding in the best set of auctions, the trader also decides to bid in other auctions which are likely to have broadly the same outcome. By this, the trader's chance of winning the desired goods is increased. Moreover, since they consider that goods have multiple attributes (for example, in auctions selling flights, the goods can be described by their dates of departure and return, by their carrier, and the class of ticket being bought), the trader has to make trade-offs between them during its bidding process in order to satisfy the user's preferences.

Furthermore, Anthony and Jennings (2002, 2006) proposed an approach for traders to buy a single good from heterogenous auctions (English, Dutch and Vickrey auctions) with varying start and end times. Their decision making model works as follows. Each trader constructs an active auction list and collects relevant information (such as current standing bid in each auction). After this, the trader determines the maximum bid it is willing to make at the current time, which is dependent on four bidding constraints: (i) the remaining time, (ii) the number of remaining auctions, (iii) the desire for a bargain and (iv) the desperation of the trader. Based on the current maximum bid, the trader determines the auctions which it can bid in and calculates what it should bid in these potential auctions. Then, the trader selects the auction with the highest expected utility as the target auction. Finally, the trader submits a bid in the target auction. In addition, they use a genetic algorithm (GA) to search for effective strategies for each of the various environments since the strategy's effectiveness is heavily affected by the environment. The trader will adopt the strategy that is most appropriate to its prevailing context by evolving strategies. In addition, Yuen et al. (2006) also investigate heuristic utility maximising bidding strategies for traders to buy one or more items from multiple heterogenous auctions with different starting and closing times. They design a two-stage strategy to approximate the best response bidding strategy for a global bidder. In the first stage, it computes a maximum bid or threshold for each auction. Then in the second stage, the trader exploits the bidding threshold calculated from the first stage and decides which auctions it should participate in.

In addition to empirical approaches to the bidding across multiple homogenous or heterogenous auctions reviewed above, Gerding et al. (2008) adopted a theoretical approach for analysing the optimal bidding strategy used in multiple, simultaneous Vickrey auctions with perfect substitutes. They consider a single bidder, called the *global bidder*, that is able to bid in any number of auctions, whereas the other bidders, called the *local bidders*, can only bid in a single auction. By theoretical analysis, they find the following results. In the multiple simultaneous Vickrey auctions context, the best strategy for a global bidder is to bid below its true private value, in contrast to that in a single Vickrey auction context, where the dominant strategy of a bidder is to bid its true private value. Furthermore, the expected utility of the global bidder is maximised by participating in all auctions even though it only requires one item. They prove that, when all auctions are identical, the strategy to maximise the global bidder's expected utility is to bid either uniformly across all auctions, or relatively higher in one of the auctions, and the same or lower in the other auctions. They argue that even though a global bidder has a considerably higher expected utility than a local one, not all bidders should necessarily bid globally. They also show that a global bidder has a positive effect on the allocative efficiency.

In addition to the above work, in the specific context of TAC Travel Competition (`http://www.sics.se/tac`), travel agents need to decide how to trade in multiple auctions in order to procure travel packages (flights, hotels and entertainment). Details about trading strategies in this context can be found in Stone et al. (2003); Vetsikas and Selman (2003); Toulis et al. (2006); He and Jennings (2003).

Now we have introduced related work about traders' strategies across multiple single-sided auctions. We can see that when buyers want to purchase multiple goods from multiple single-side auctions, they first need to decide which auction(s) they want to participate in, and then decide what bids to submit. This analysis will be insightful for us to analyse how traders select marketplaces and submit shouts in the context of multiple competing double auction marketplaces. However, our analysis in the double auctions will be much more complex. The reason is as follows. In multiple single-side auctions where there is a seller staying in each single-sided auction, researchers only need to consider the market selection of buyers. However, in the double auctions, there are multiple buyers and multiple sellers. We need to consider the market selection of both buyers and sellers. The market selection of buyers(sellers) not only depends on market policies and market fees, but also depends on the number and types of other traders in marketplaces. Furthermore, in single-sided auctions, only the buyer bidding the highest wins the good, and thus how buyers bidding depends on other buyers' bidding behaviour. However, in double auctions, the transactions can happen between multiple buyers and multiple sellers. Thus how buyers(sellers) submitting shouts not only depends how other buyers(sellers) bidding, but also depends on how sellers(buyers) bidding.

2.4.1.2 Competition between Single-Sided Marketplaces

So far we have described related work about traders' strategies in the context of multiple single-sided auctions. However, this work does not consider the interaction between auctions themselves. Actually, auctioneers (or sellers, when sellers determine the auction rules) need to compete with each other to attract traders (in terms of, e.g. charging fees, setting the duration of the auction, setting reserve prices, and so on). In the following, we will describe the related work about how single-sided auctions compete against each other.

McAfee (1993) was the first to consider mechanism design and reserve prices (which is the lowest price that the seller will agree to sell the good for) in the context of competing sellers. In his paper, sellers can choose any direct mechanism and these mechanisms are conducted for multiple periods with discounted future payoffs. He studies the relationship between the auction mechanism (which is determined by the sellers) and type distributions of buyers across sellers, and the consequences of this relationship for sellers choosing an auction mechanism in equilibrium. He finds that the equilibrium reserve prices posted by sellers are equal to the sellers' values. However, his work is based on some strong assumptions (e.g. it assumes that any individual seller has no significant impact on buyers' profits, and that buyers' expected profits in future periods are invariant to deviation of a seller in the current period), which are only reasonable in the case of infinitely many players. Burguet and Sakovics (1999) then relax some of these strong assumptions. Specifically, they derive a unique equilibrium strategy for the buyers when two sellers are competing with each other by setting reserve prices. They prove that there always exists an equilibrium strategy for the sellers, but it cannot be a symmetric one in pure strategies. They show that in the mixed equilibrium, sellers will set a reserve price above their own private value of the good. Another work about the competition between sellers is Hernando-Veciana (2005), which considers a large number of sellers. This paper shows that each seller announces a reserve price equal to its private values when the number of sellers and buyers is large.

Another important work is done by Gerding et al. (2007), which considers the setting with a small numbers of buyers and competing sellers, and considers that a mediator charges fees to the competing seller for running the auction. This work shows that pure Nash equilibria for the asymmetric seller setting exist. In addition, they also find that by shill bidding (where the seller pretends to be a buyer to bid in its own auction), the seller can further improve its utility. Finally, they evaluate the ability of different fees to deter shill bidding and quantify their impact on market efficiency. They show that auction fees based on the difference between the payment and the reserve price are more effective than the more commonly used auction fees with regards to deterring shill bidding and increasing market efficiency.

Now we have introduced the related work about the competition between single-sided auctions. In this competition, sellers who determine the auction rules, compete with each other to attract buyers by setting reserve prices. Therefore, this work only considered single-sided auctions attracting one side of traders. Our double auctions will be more complex since they have to compete with each other to attract both sides of traders (i.e. buyers and sellers), and the market selection strategies of buyers and sellers also interact with one another.

2.4.2 Two-Sided Marketplaces

After describing related work about competing single-sided auctions, we now turn to two-sided double auction marketplaces with multiple buyers and multiple sellers. So far there exists no literature that deals with trading strategies which are specifically designed to operate across multiple double auctions (in Section 4.3, we will use fictitious play to analyse traders' equilibrium strategies across multiple marketplaces). However, for the setting with single-home trading where traders can only enter one marketplace at a time, the bidding strategies introduced in Section 2.3.2.1 (i.e. ZI-C, ZIP, GD, RE, TT, PvT, GDX, AA, FL and so on), which are designed for isolated marketplace environment, are appropriate and efficient for traders to use. Then to select which marketplace to choose, Niu et al. (2008b) propose two market selection strategies for traders to search for the most profitable marketplace (we will introduce them in Section 2.5.1 when we describe the CAT competition). Note that there is comparatively little theoretical work focusing on the competition between double auction marketplaces. However, there exists a considerable body of work on analysing competing two-sided marketplaces. As we mentioned in Section 1.1, the double auction marketplace is a particular type of two-sided marketplace. Thus research on general two-sided marketplace competition should benefit research on the more specific double auction marketplace setting. In the following, we describe the existing work about how two-sided marketplaces compete with each other, where researchers mainly focus on determining fees charged to traders (i.e. charging strategy).

2.4.2.1 Competition between Two-Sided Marketplaces

There exist various two-sided marketplaces in the real world. Examples of such two-sided marketplaces are dating websites where males are on one side and females are on the other side, operating systems for PCs where software developers are on one side, and consumers are on the other side, card payment systems where shops are on one side and cardholders are on the other side, and so on.

In a two-sided marketplace, the traders' choice of marketplaces is significantly affected by the number of traders on the other side. This comes from the *positive size effect* (Caillaud and Jullien, 2003), which means that traders have larger expected profits in the marketplace which has the larger number of traders on the other side. This is because a large number of one side gives the other side access to more diversity. Thus, the two-sided marketplaces need to try to attract traders on both sides (since a small number of traders on one side will cause less expected profits for traders on the other side and, in turn, cause those traders to leave the marketplace). For example, payment card systems need to attract both retailers and cardholders to adopt their cards, operating systems need to court both users and application developers, and dating websites need to attract both males and females. However, this gives rise to a “chicken and egg” problem: to attract buyers (one side), a marketplace should have a large base of the registered sellers (the other side), but these will be willing to register only if they expect many buyers to show up.

One of the most important works on this problem is from Caillaud and Jullien (2003), who analyse the competition between two marketplaces. In their work, they assume that traders are homogeneous with the same preference (type), and traders’ utilities only depend on the number of traders on the other side. Furthermore, two types of fees are allowed in their model: a registration fee charged when traders enter the marketplace, and a transaction fee charged when traders are successfully matched. In addition, they allow the marketplaces to charge negative registration fees, which means that the marketplaces can subsidize traders in order to attract them. Firstly, they analyse the case with single-home trading. Specifically, they analyse a “divide-and-conquer” strategy of subsidizing one side of the traders (by charging negative registration fees) while recovering loss (by charging positive transaction fees) from the other side. They show that, when traders can only enter one marketplace at a time, competition between marketplaces is severe and, in equilibrium, all traders enter the same marketplace, but this marketplace has to give up all profit. They further analyse the more complicated case with multi-home trading, which induces lower degree of competition and, as a result, generates positive profits in equilibrium. They show that in the multi-home environment, an efficient equilibrium always exists in which traders register with the marketplace with cheaper registration fees, and are willing to register with the marketplace which provides the lower transaction fee.

Another significant work on competition in two-sided marketplaces is Rochet and Tirole (2003), whose analysis can be understood in the context of competing credit card platforms with two sides: consumers and the set of retailers. In their analysis, the credit card platforms charge fees for each transaction. When a credit card offers a lower transaction fee to retailers than its opponent, then a retailer needs to make a choice between only accepting the cheaper card or accepting both cards. When the retailer accepts only the cheaper card, then its consumers have a stark choice between paying by this card or not using this card at all. However, if the retailer accepts both cards, then it may happen that fewer consumers use the retailer’s preferred lower-cost card to pay. By theoretical analysis they find that, in equilibrium, all retailers accept both credit cards while consumers use their preferred cards. The split of transaction fee charged to both sides depends on how consumers consider the substitute of the two competing cards. If they consider the two cards as close substitutes, then retailers need to bear most of the transaction fee in equilibrium. They also derive a simple formula which can be used to govern fees in two-sided marketplaces. However, in both Caillaud and Jullien (2003) and Rochet and Tirole (2003), it is assumed that buyers(sellers) are homogeneous with the same preference (type), and will choose the same marketplace. In the real-world double auctions, this assumption is not appropriate since traders are usually heterogeneous with different preferences. In addition, Damiano and Hao (2008) consider a setting with heterogeneous traders. They assume that traders select marketplaces only according to the types of other traders, instead of the number of traders, and thus they do not consider the (positive) size effect.

Furthermore, another important work is done by Lee (2008), who analyses under what conditions two competing marketplaces can co-exist. They point out that strong marketplace differentiation or weak positive size effect can make competing marketplaces co-exist. In our work, we will also analyse whether competing marketplaces can co-exist and under what conditions if they can.

Although the above works are related and can benefit the research on competing double auction marketplaces, none of these papers specifically consider the double auction mechanism to match traders. This changes the problem because in the competition of double auction marketplaces, the market selection strategy of the traders not only depends on the number of traders choosing the marketplaces, but also on their own types and those of other traders choosing the marketplaces. Furthermore, in addition to the positive size effect, there also exists the *negative size effect*. Buyers(sellers) have to compete with each other in order to be matched with sellers(buyers). Therefore, buyers(sellers) prefer the marketplace with few buyers(sellers). Competition between two double auction marketplaces is considered by Ellison et al. (2004). They show that, in some cases, the negative size effect has a larger impact than the positive size effect, and traders will not migrate from this state, which means that two competing marketplaces can co-exist in equilibrium. This model is similar to ours since we also consider heterogeneous traders and both positive and negative size effects. However, Ellison et al. (2004) make the simplifying assumption that traders choose a marketplace before learning their own types, and thus the market selection strategy is independent of a trader's type. Therefore, unlike in our model, using their model, they show that similar marketplaces can co-exist in equilibrium. In contrast, we find that in our model, traders will converge to one marketplace except for the case when there is strong market differentiation (i.e. where competing marketplaces charge different types of fees).

2.5 The CAT Competition

After describing the related work about competing marketplaces, now we introduce a specific competition between double auctions since our design of a competing marketplace will be evaluated in this context. In order to promote the research on competing double auction marketplaces, an annual Market Design Competition (CAT) was introduced as part of the Trading Agent Competition (TAC) (Cai et al., 2009). To this end, a marketplace competition platform was provided, which is called JCAT (Niu et al., 2008b). In the following, we will detail this platform according to technical reports issued by CAT competition organisers (see Cai et al. 2009; Niu et al. 2009) and describe the related work in this specific context of CAT.

2.5.1 Basic Structure and Rules

The CAT competition was first introduced as part of TAC in 2007, in order to promote research on efficient and effective competing marketplace design. The underlying platform, JCAT, allows multiple marketplaces to compete against each other and allows marketplaces to be evaluated in a uniform way. Furthermore, this competition has been run successfully in 2007, 2008, 2009 and 2010, and can provide an international benchmark for evaluating the competing marketplace design. Given this, in this thesis, we will evaluate our design of a competing marketplace based on this platform.

As depicted in Figure 2.10, the CAT competition consists of traders, i.e. buyers and sellers, and specialists. Traders can buy or sell goods in one of the available marketplaces which are operated by specialists. In the competition, the traders are provided by the CAT organisers, and specialists (with the market policies and charging strategies) are designed by the competition entrants. Each entrant can only operate a single marketplace in the competition.

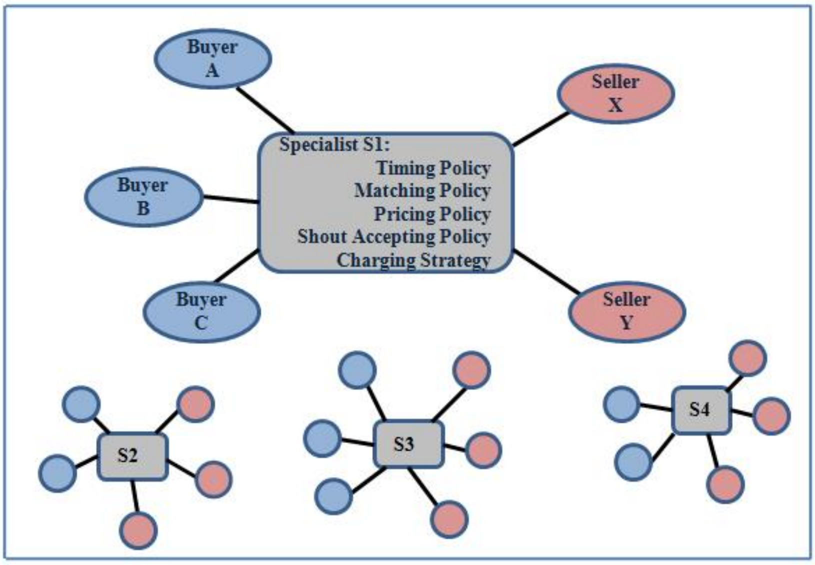

FIGURE 2.10: Architecture of CAT competition.

Each CAT competition lasts for a number of trading days[7], and each day consists of a fixed number of rounds, during which traders submit shouts to the specialist they are registered with. Each round lasts for a known constant length of time. In the competition, each trader is assigned a private value for the goods it wishes to buy or sell. For the buyer, the private value is its limit price, and for the seller, the private value is its cost price. In the competition, traders' private values remain constant during a day, but may change from day to day, depending on the configuration of the game parameters. Traders need to enter marketplaces to make transactions. Here, each trader can only register with one marketplace on a particular day, i.e. only single-home

[7]A trading day is a virtual day, not an actual tournament day in real time.

trading is considered. Thus bidding strategies in isolated marketplaces can be used. Typically, ZI-C, ZIP, GD and RE strategies (see Section 2.3.2.1 for details) are permitted in this setting. In addition to the bidding strategies, traders have one of several market selection strategies to decide which marketplace to register with. These market selection strategies implemented in JCAT are as follows:

- *random strategy*: the trader randomly chooses a marketplace to participate in.
- *ϵ-greedy exploration strategy*: the trader treats the choice of marketplace as an n-armed bandit problem[8] which it solves using an ϵ-greedy exploration strategy where $0 \leq \epsilon \leq 1$. With such a strategy, a trader updates its value function according to its recent profits, and then it chooses the most profitable marketplace with probability $1 - \epsilon$, and randomly chooses one of the remaining marketplaces with probability ϵ. The probability ϵ can remain constant or can vary over time, depending on the value of a parameter α. If α is 1, ϵ remains constant, while if α takes any value in $(0, 1)$, then ϵ will decrease over time.
- *softmax exploration strategy*: This market selection strategy is similar to the above ϵ-greedy exploration strategy except that it uses a softmax exploration strategy in the n-armed bandit algorithm. This means the trader does not treat all marketplaces, except the best marketplace, as the same. If this strategy does not choose the most profitable marketplace, it weighs the choice of remaining marketplaces in order to choose more profitable marketplaces. A parameter τ, which is similar to ϵ, controls the relative importance of the weights a trader assigns to marketplaces. It can also be fixed or have a variable value determined by α.

Specialists facilitate transactions by matching bids and asks and determining the transaction price. As discussed in Section 1.1, a double auction marketplace consists of the following market policies: *timing*, *matching*, *pricing* and *shout accepting policies*. Furthermore, the marketplace also needs to specify its *charging strategy*. An overview of the policies implemented in JCAT is listed in Table 2.4. In the pricing policy, in addition to k-pricing policy and n-pricing policy we introduced in Section 2.3.2.2, JCAT implements two other pricing policies:

- *Uniform pricing policy*: sets the transaction prices for all matched ask-bid pairs at the same point.
- *Side-biased pricing policy*: is a kind of k-pricing policy with varying k, where k is set to split the profit in favour of the side where fewer shouts exist.

In addition to the quote-beating accepting and equilibrium accepting policy introduced in Section 2.3.2.2, JCAT implemented four other accepting policies:

[8] This is a machine learning problem, similar to the slot machine, where there are n levers, and each lever provides a reward drawn from a distribution associated with that lever Sutton and Barto (1998). The objective is to maximize the reward sum through iterative pulls. There is no initial knowledge about the levers. In each pull, a trade-off is made between "exploitation" of the lever that has the highest expected reward and "exploration" to get more information about the expected rewards of the other levers.

- *Always accepting policy*: accepts any shouts submitted by traders.
- *Transaction-based accepting policy*: records the most recently matched bids and asks, and then uses the lowest matched bid and the highest matched ask as thresholds to restrict the allowed shouts.
- *History-based accepting policy*: is derived from the GD bidding strategy. GD forms the belief about how likely a given shout is to be matched based on the history of previous shouts. Then, according to this belief, the marketplace accepts shouts that can be matched with the probability no lower than a specified threshold.
- *Self-beating accepting policy*: accepts all shouts of traders who have not submitted shouts yet, but then only allows traders to modify their standing shouts with more competitive prices.

For the charging strategy, JCAT provides five types of fees which can be set:

- **Registration fee**: a flat fee charged for participating in a marketplace.
- **Information fee**: a flat fee charged by the specialist to traders and specialists who require market information from the marketplace. This information consists of the entire history of the accepted shouts and the transactions in that marketplace.
- **Shout fee**: a flat fee charged for successfully submitted bids or asks.
- **Transaction fee**: a flat fee charged on each successful transaction.
- **Profit fee**: a share of the observed profit made by traders, where a trader's observed profit is calculated as the difference between the shout and transaction price. Note that a trader's observed profit is different from its actual profit (which is the difference between its private value and transaction price).

Given this, JCAT comes with five pre-set charging strategies. The five pre-set charging strategies in the JCAT platform are as follows:

- *Fixed charging*: imposes fees at a specified fixed level.
- *Bait-and-switch charging*: makes a specialist reduce its charges until it captures a certain market share, and then it slowly increases charges to increase profit. The specialist adjusts its charges downward again when the market share drops below a certain level.
- *Charge-cutting charging*: sets the charges to the lowest charges set by the other marketplaces on the previous day. This will attract many traders, but may sacrifice profits.
- *Learn-or-lure-fast charging*: adapts charges towards some desired target according to the scheme used by the ZIP bidding strategy.

Timing Policy (Section 2.3.2.2)	**Matching Policy** (Section 2.3.2.2)
Continuous clearing[9] Round clearing[10]	ME
Pricing Policy (Section 2.3.2.2)	**Charging Strategy**
k-pricing *n*-pricing Uniform pricing Side-biased pricing	Fixed charging Bait-and-switch charging Charging-cutting charging Learn-or-lure-fast charging

Shout Accepting Policy (Section 2.3.2.2)
Quote-beating accepting Estimated equilibrium Accepting Self-beating accepting Transaction-based accepting History-based accepting

TABLE 2.4: Market policies implemented in the JCAT.

Although JCAT has provided some pre-set market policies and charging strategies, specialists can design *their own* in order to effectively compete with one another in making profit, attracting traders and ensuring shouts submitted in the marketplace result in transactions. To evaluate their performance, the marketplaces are scored based on a combination of three different metrics:

- **Profit Share**: the profit obtained by the specialist on a particular day as a percentage of the total profits obtained by all specialists on that same day. The profit share score is a number between 0 and 1 for each specialist for each day.
- **Market Share**: the percentage of traders who have registered with that specialist on a given day. The market share score is between 0 and 1.
- **Transaction Success Rate**: the percentage of shouts accepted by the marketplace resulting in transactions. The transaction success rate score is between 0 and 1.

These metrics are weighted equally (i.e. weighted one-third each) and added together to produce a combined score for each specialist for each Assessment Day[11]. Scores are then summed across all Assessment Days to produce a final game score for each specialist. The specialist with the highest final score is the winner of the competition.

2.5.2 Traders Migrating between Marketplaces

In the competing marketplace context with single-home trading, traders move freely between the different marketplaces to search for the one which they believe to be the most profitable.

[9]It is the implementation of CDA policy in the JCAT.
[10]It is the implementation of CH policy in the JCAT.
[11]In order to avoid effects arising from the fact that the competition has a start day and an end day, not all the trading days will be used for assessment purpose. In more detail, the organisers randomly choose a day as a starting trading day and a day as an ending day, and then randomly choose days between the starting day and the ending day on which assessment will be undertaken. These days are called "Assessment Days".

Such migration will significantly determine the final competition results of marketplaces, and thus needs to be researched in detail. To this end, based on the JCAT platform, Niu et al. (2007) provided an experimental analysis on this aspect of the traders' behaviour. Specifically, in their experiments, traders' market selection strategies are chosen from random selection strategy, ϵ-greedy exploration strategy or softmax exploration strategy, which we described in Section 2.5.1. From the results, the authors concluded that, when traders choose marketplaces randomly, marketplaces charging higher fees tend to make larger profits. However, as soon as traders are able to learn, they exhibit a strong tendency to concentrate towards marketplaces charging low fees. They also showed that the traders' market selection is dependent on *how* traders learn. In particular, using ϵ-greedy exploration allows traders to settle on the cheaper marketplace more quickly, while the softmax exploration allows traders to better distinguish the range of non-optimal marketplaces.

Building on this, Cai et al. (2008) ran experiments to analyse the impact of multiple marketplaces on the efficiency of trading. From their results, they found that, while dividing traders into multiple marketplaces leads to a loss of allocative efficiency, this loss is reduced when traders are allowed to move among marketplaces in search of greater profits. In particular, they showed that when traders can move among marketplaces, this can cause segregation. For example, most intra-marginal traders may concentrate on one marketplace because, in this marketplace, they have more possibilities to trade. In the specific context of CDAs, extra-marginal traders may, in general, be able to steal a deal from intra-marginal traders, but when such segregation occurs, it is hard for extra-marginal traders to do so. This, in turn, leads to increased allocative efficiency. In addition, they showed that fees can drive traders that are not making profits to try different marketplaces, and thus allow marketplaces to rid themselves of unproductive traders. In the CHs, although extra-marginal traders cannot steal a deal from intra-marginal traders, the movement of traders can still increase profits by allowing a trader that is extra-marginal in one marketplace to become intra-marginal in another. Generally speaking, because traders are profit-motivated, they migrate towards marketplaces that provide more profits, and overall this increases the total profits of the set of marketplaces, increasing the global efficiency. This effect is enhanced by the application of fees since these tend to reduce profits and then discourage traders from remaining in marketplaces that are unprofitable for them.

To sum up, from the above experimental analysis, we can see that profit-motivated traders with a learning ability will migrate towards profitable marketplaces, and this migration process is enhanced by fees imposed by marketplaces since they reduce traders' profits. The way that traders move between marketplaces is important when considering the design of market policies. Effective market policies are needed to provide highly efficient allocation in order to guarantee traders' profits and attract traders. Moreover, the charging strategy should be given particular consideration since it has been shown to have a large impact on the traders' decision of which marketplace they will register with. However, this experimental analysis considered heuristic strategies adopted by traders and marketplaces and the results depend on the choices of heuristics, and thus cannot easily be used to derive general conclusions.

In addition, in Sohn et al. (2009), PSUCAT proposes a simple game-theoretic model to analyse traders' market selection in the CAT context, which assumes a game with complete information about the traders' private values. They analyse how registration fees affect the market selection of intra-marginal and extra-marginal traders. They show that extra-marginal traders are the first to leave the marketplace when it starts charging a registration fee. Intra-marginal traders are the next to leave the marketplace as the registration fee increases. Furthermore, they find that the resulting number of Nash equilibria drop as the registration fee increases. However, this analysis assumes a game with complete information about the traders' types, in contrast to reality, where traders' types are usually privately known.

2.5.3 Competition Results

In this section, we introduce the market policies and charging strategy used by specialists in the CAT competition (2007, 2008, 2009 and 2010), where IAMwildCAT is the entry of University of Southampton, that I have been involved since 2008. This existing work, which is directly related to the competing marketplace design in the context of the CAT competition, will give us further understanding of the design of competing marketplaces.

2.5.3.1 2007

The 2007 CAT Competition consisted of 10 specialists: IAMwildCAT, PSUCAT, CrocodileAgent, jackaroo, Havana, PersianCat, PhantAgent, Mertacor, TacTex and Manx, and the result is shown in Table 2.5. In Niu et al. (2010), the competition organisers analyse the policies used in the competition by inferring the market policies from the binaries of the specialists submitted by the participants[12]. They analysed the competition through two approaches: while-box approach and black-box approach. The white-box approach attempts to relate the internal logic and features of strategies to the competition outcomes. In more detail, they showed that most specialists use ME to clear markets at the equilibrium price. IAMwildCAT and Mertacor are the only two which attempt to match intra-marginal shouts with extra-marginal shouts close to the equilibrium point in order to reach high transaction success rates. Furthermore, about half the specialists use a round clearing policy which means the marketplace is cleared at the end of each round, and the other half use a continuous clearing policy. The specialists in the competition used a wide range of shout accepting policies. In particular, IAMwildCAT, PSUCAT, CrocodileAgent and Mertacor used a modified or improved estimated equilibrium shout accepting policy. Havana, jackaroo and MANX adopted the quote-beating accepting policy. TacTex used an always-accepting policy and PersianCat used a transaction-based accepting policy. For the pricing policy, Havana, PersianCat, TacTex and MANX used the *k*-pricing policy. On the other hand, IAMwildCAT, PSUCAT and Mertacor used a modified or improved side-biased pricing policy. jackaroo used

[12]The CAT competition organisers require the participants to submit the binaries of the specialist they used in the competition, but not the actual source code.

Rank	Specialist	Score
1	IAMwildCAT	365.496
2	PSUCAT	328.333
3	CrocodileAgent	308.813
4	jackaroo	272.239
5	Havana	256.441
6	PersianCat	218.458
7	PhantAgent	146.208
8	Mertacor	133.912
9	TacTex	107.925
10	MANX	95.976

TABLE 2.5: Result of 2007 CAT competition

the n-pricing policy, whereas CrocodileAgent used a combination of modified n-pricing policy and side-biased pricing policy. From all the market policies, most effort seems to have been placed in the charging strategy and most specialists designed their own or modified charging strategies provided by JCAT. Considering how fees are updated over time, some specialists adapted their charges, whereas other specialists directly calculated the charges such that they would expect to achieve a certain target profit, and some specialists combine the above two approaches by changing the fees gradually from the current level to the target level. For different type of fees, about half of the specialists mainly or exclusively charged registration fees and profit fees. TacTex charged only shout fees. Three specialists, CrocodileAgent, Havana and MANX charged each kind of fees without any preference on a certain kind of fee. Moreover, most specialists utilise the *start effect* (which means attracting traders by charging less in the early stage of game) and the *end effect* (which means imposing higher charges to earn more profit when the competition is about to end) even though the starting and ending days may not be included in the Assessment Days.

Furthermore, they use a black-box approach to analyse the competition, which will consider the strategies as atomic entities. A black-box analysis abstracts away the internal structure of the specialists and many details of the dynamics during the interaction between specialists, making it possible to consider many more situations. In more detail, in this approach, they consider two sets of experiments: multilateral simulations with games involving all the specialists and bilateral simulations with games each involving two specialists. Through the analysis, they show that in these simulations, the winner IAMwildCAT still dominates other entrants. They further introduce a specialist MetroCat (which is designed by organisers for evaluation purpose), and show that this specialist quickly dominates all entrants including IAMwildCAT. The success of MetroCat also suggests that there is a significant room for the improvement of entrants in this year's competition.

Rank	Specialist	Score
1	PersianCat	425.773
2	MANX	384.921
3	jackaroo	380.554
4	PSUCAT	379.416
5	Mertacor	354.283
6	IAMwildCAT	331.708
7	DOG	314.014
8	CrocodileAgent	287.8
9	MyFuzzy	259.503
10	BazarganZebel	252.128
11	Hairball	242.853

TABLE 2.6: Result of 2008 CAT competition

2.5.3.2 2008

In the 2008 CAT Competition, there were 11 specialists: BazarganZebel, CrocodileAgent, DOG, Hairball, IAMwildCAT, MANX, Mertacor, MyFuzzy, PSUCAT, PersianCat and jackaroo, and the result is shown in Table 2.6. Up to now, only three entrants have published their strategies in the 2008 CAT Competition: PersianCat (the winner of the 2008 CAT Competition), Mertacor and CrocodileAgent. In the following, we introduce their policies briefly, and then introduce the analysis of the 2008 CAT Competition by the competition organisers.

In the 2008 CAT Competition, the market policies adopted by PersianCat are as follows (Honari et al., 2009):

- *timing policy:* it adopts continuous clearing policy, which clears the marketplace immediately when a new shout is admitted.
- *matching policy:* it uses the equilibrium matching policy, i.e. matching high bids with low asks.
- *shout accepting policy:* it adopts an equilibrium accepting policy. Specifically, it calculates the equilibrium price on day i according to the equation: $\hat{p}_i^* = \frac{1}{2n}\sum_{j=i-1}^{j=i-n}(MaxA_j + MinB_j)$, where n is the length of sliding window, and is equal to 4 in the competition, $MaxA_j$ and $MinB_j$ are the maximum transacted ask and the minimum transacted bid on day j. Then it set a slack value ϵ to moderate the restriction of the accepting policy. It sets $\hat{p}_i^* - \epsilon$ and $\hat{p}_i^* + \epsilon$ as the minimum acceptable bid and the maximum acceptable ask respectively in PersianCat on day i. In the competition, the ϵ is 10% of the $\hat{p}_i^*$.
- *pricing policy:* it sets the transaction price to its estimated equilibrium price. If the equilibrium price is outside the range of the matched bid and ask, it sets the transaction price to the nearest price of the matched bid and ask.
- *charging strategy:* it adopts a charging strategy, which only charges a fixed profit fee.

The market policies adopted by Mertacor are as follows (Stavrogiannis and Mitkas, 2009):

- *timing policy:* it uses a round clearing policy for the earlier rounds, and then switches to a continuous clearing policy to increase the volume of its transactions.
- *matching policy:* it uses the equilibrium matching policy, i.e. matching high bids with low asks.
- *shout accepting policy:* it designs a global equilibrium accepting policy, which allows only globally intra-marginal traders to place shouts in the earlier rounds of a given trading day. It then subsequently switches to the quote-beating accepting policy. Specifically, the global equilibrium price is calculated as follows. It continually keeps track of the highest bids and the lowest asks in Mertacor. It believes that these prices constitute the closest available estimation of traders' private values. Moreover, it estimates the number of goods traded every day as traders' daily endowment. When a sufficient number of traders have been explored, it forms the global demand and supply curves and then computes the global equilibrium price.
- *pricing policy:* it adopts a uniform global equilibrium pricing policy, which sets the price of all transactions at the global equilibrium price.
- *charging strategy:* it only charges a profit fee, and sets the fee according to its market statistic and opponent scores.

The market policies adopted by CrocodileAgent are as follows (Petric et al., 2008):

- *timing policy:* it clears the market every second round.
- *matching policy:* it uses the equilibrium matching policy, i.e. matching high bids with low asks.
- *shout accepting policy:* it uses an equilibrium accepting policy provided by the JCAT platform.
- *pricing policy:* it adopts a slightly modified n-pricing policy in order to reduce the loss of traders and ensure that both buyers and sellers will obtain profits from transactions.
- *charging strategy:* it divides the game into two phases. In the luring phase, it tries to attract traders by charging no fees. Then once obtaining a good market share, it switches to the charging phase, and begins to charge fees to make profits.

The competition organisers also analysed specialists in the 2008 CAT Competition in order to better understand the characteristics of the policies adopted by specialists in relation to the competition context (Robinson et al., 2009). Specifically, they empirically evaluate the generalisation abilities of marketplaces in the different trader populations, different competing specialists populations and different scoring periods. They show that specialists in the 2008 CAT

Rank	Specialist	Score
1	jackaroo	542.6
2	CUNY.CS	533.1
3	IAMwildCAT	512.9
4	PSUCAT	499.7
5	Mertacor	395.0
6	rucat0	391.2
7	UMTac09	379.4
8	cestlavie	365.1
9	PersianCat	327.0
10	BazarganZebel	315.6
11	TWBB	294.0
12	Tianuani	261.0
13	CrocodileAgent	217.7
14	WaterCAT	167.1

TABLE 2.7: Result of 2009 CAT competition

Competition are not robust, which means that a specialist which performs well in a particular competition context may not perform well if the competition context changes. Thus the market policies adopted by specialists seem not to generalise, and the competition organisers suggest entrants should make their specialists more robust in the future competition.

2.5.3.3 2009 and 2010

The results of the CAT competition in 2009 and 2010 are shown in Tables 2.7 and 2.8 respectively. We found that jackaroo performed well in these two years' competition (ranked first in 2009 and second in 2010), and Mertacor's performance was improved significantly from the fifth in 2009 to the first in 2010. We also can see that our entrant, IAMwildCAT, performed well in these two years' competition (ranked third in both years). Furthermore, we found that the performance of competition marketplaces changed in different days' competition since competition operators changed traders' strategies. Up to now, neither entrants nor competition organisers have published their analysis for these two years' competition. Therefore, we do not know market policies and charging policies used by other competing marketplaces. However, for the market fee charged to traders, which is the information we can directly observe from the competition, we find that most effective marketplaces only charge profit fees, and set other types of fees as zero.

Rank	Specialist	Score
1	Mertacor	627.394
2	jackaroo	589.667
3	IAMwildCAT	548.238
4	PoleCAT	521.918
5	TWBB	461.398
6	PSUCAT	447.459
7	MyFuzzy	428.547
8	AstonCAT	393.911
9	PersianCat	255.724

TABLE 2.8: Result of 2010 CAT competition

2.5.3.4 IAMwildCAT

In this section, we introduce the market policies and charging policy used by IAMwildCAT in the 2007, 2008, 2009 and 2010 CAT competition.

2007:

In terms of the timing policy, IAMwildCAT clears the market at the end of each round, i.e. round clearing. In terms of the matching policy, it is observed that intra-marginal traders are expected to trade earlier than marginal traders such that the amount of profit to be extracted in the marketplace is higher earlier in the trading day, and with less profit to be made at the end of trading day. Thus IAMwildCAT chooses the following strategy to deal with the trade-off: it clears the market for maximising profits in the end of earlier rounds of the trading day. Then, on the following rounds, with less profits to be made in the marketplace, IAMwildCAT clears to maximise the number of transactions.

For the pricing policy, IAMwildCAT uses a variation of the *k*-pricing policy. In particular, IAMwildCAT looks at a window of the 10 trading days for the average number of buyers and sellers it attracts. If the difference between the number of buyers and sellers is bigger than 10% of the total number of traders, *k* is adjusted in order to give more profit to the side (buyers or sellers) which is under represented. Such a pricing policy is similar to the side-biased pricing policy implemented in the JCAT.

In the design of the shout accepting policy, IAMwildCAT adopts the equilibrium shout accepting policy to reject bids below the equilibrium price and asks above the equilibrium price. Because of the error of estimating the equilibrium price, IAMwildCAT provides some slack when determining the minimum accepted bid and the maximum accepted ask. Furthermore, in order to increase TSR, on the last few trading rounds, IAMwildCAT only accepts bids and asks that can be currently cleared.

Finally, in the design of the charging policy, initially, IAMwildCAT explores the marketplace in order to build up its market share. Specifically, it attracts as many traders as possible by giving up all profits, scoring only by its higher market share and its transaction success rate. With a large market share, it then starts exploiting these traders by charging registration and profit fees to extract profits from them. With increasing fees, the marketplace becomes less attractive to potential traders, and then gradually, the effect of increased fees decreases the market share. At this point, IAMwildCAT stops exploiting and goes back to exploring to increase its market share back to a predetermined threshold. Once it reaches that target, it has a sufficient number of traders registered to start exploiting. By following this policy, the marketplace can maintain the market share at a reasonably high level while exploiting the traders whenever the market share is very high. This charging policy would then oscillate between the exploiting and exploring behaviours until the end of the game.

2008:

In 2008, in the design of IAMwildCAT, for the timing policy, the same as 2007, the marketplace adopts the round clearing policy. For the matching policy, IAMwildCAT makes a trade-off between maximising traders' profits and maximising the number of transactions. In more detail, when shouts are far away from the equilibrium price, it uses the ME matching policy to maximise traders' profits. Then when traders' shouts are within an area close to the equilibrium price, it uses the MV matching policy to maximise the number of transactions. For the pricing policy, it is the same as that in 2007 CAT competition. For the shout accepting policy, IAMwildCAT uses a quote-beating accepting policy. Finally, for the charging policy, it adopts a concept of positive feedback loop. In more detail, IAMwildCAT begins to charge fees until it obtains the highest market share, and because of the highest market share, it may obtain the highest market profit among all competing marketplace by charging lower fees. Because of lower fees, the marketplace can still maintain and attract more traders. Then attracting more traders means charging even lower fees to traders to obtain the highest market profit. In this situation, we can see that IAMwildCAT can charge lower fees to obtain good market profit, but still remain competitive.

2009:

In 2009, IAMwildCAT uses the same timing and matching policies as those in 2008. For the pricing policy, it uses an equilibrium pricing policy, where the equilibrium price is set as the transaction prices. The estimation of the equilibrium price is similar to that PersianCat did in 2008 (see Section 2.5.3.2). For the shout accepting policy, IAMwildCAT uses an equilibrium-beating accepting policy. For the charging policy, from the competition of 2008, it is observed that IAMwildCAT fails to obtain the highest market share, but is able to keep market share at a good level. In this situation, the charging policy used in 2008 does not charge enough to make profit. In this year, the charging policy is improved according to this observation. In more detail, when the marketplace fails to obtain the highest market share, but can keep it at a good level, the marketplace still begins to charge fees.

2010:

In 2010, according to insights from the analysis done in this thesis, we design market policies and the charging policy. In more detail, IAMwildCAT makes a trade-off between the round clearing and the continuous clearing policies. It also makes a trade-off between maximising the number of transactions and maximising traders' profits. For the pricing policy, it adopts the equilibrium pricing policy, and for the shout accepting policy, IAMwildCAT switches between the quote-beating and the equilibrium beating accepting policies. For the charging policy, IAMwildCAT uses an adaptive charging policy which adapts fees based the number of transactions. For details, refer to Chapter 4.

2.6 Summary

In this chapter, we began by outlining the background on game theory and market theory in microeconomics. We then introduced related work on isolated single-sided auctions and double auctions. After this, we introduced existing work about competing single-sided auctions and competing two-sided marketplaces. Finally, we introduced the CAT competition for evaluating the competing double auction marketplace and reviewed the relevant work in this specific context.

Although existing work analyses the competition between marketplaces, they do not fully address the research challenges introduced in Section 1.2. In the following, we discuss why they fail to address our research challenges in detail, and briefly introduce what we do in this thesis to address these challenges.

Firstly, from the literature review, we can see that a number of theoretical models have been proposed to analyse how traders select two-sided marketplaces and how competing two-sided marketplaces set fees to make profits and keep traders. They are related to our work on competing double auction marketplaces since the double auction marketplace is a particular type of a two-sided marketplace. However, most of this work only considers the *positive size effect*, whereby buyers(sellers) prefer marketplaces which have a larger number of sellers(buyers). As we discussed in Section 1.2, in double auction marketplaces, the negative size effect also exists, since traders on one side will compete with each other in order to be matched with traders on the other side. This negative size effect will encourages traders to distribute across different marketplaces, thereby making it more likely for several competing marketplaces to co-exist in the long term. In this thesis, our analysis of competing double auction marketplaces will inherently consider both positive and negative size effects, and we find that the positive size effect has a larger impact than the negative size effect, and traders will concentrate in one marketplace (see Chapters 3 and 4).

Moreover, most existing work assumes that all traders are homogeneous (i.e. have the same type), and the marketplaces have complete information about the types of traders. These simplifying assumptions fail to fulfill our research challenge that traders usually have heterogeneous types which are privately known. In this situation, the traders, in choosing their marketplaces, not only care about fees charged by the marketplaces and the number of other traders, but also their types. Although in the CAT competition there exists some work that considers these factors, they consider specific heuristic strategies adopted by traders and marketplaces and the results depend on the choices of heuristics, and thus cannot easily be used to derive general conclusions. In our analysis, we will consider that traders have privately known heterogeneous types. Specifically, in Chapter 3, we analyse equilibrium strategies of traders and marketplaces under the assumption that traders have heterogeneous discrete trader types which are privately know. In Chapter 4, we extend the analysis to the setting with privately know continuous trader types.

Furthermore, comparatively little existing work considers different trading environments (single-home, multi-home and hybrid trading) and currently no literature considers different good properties (independent, substitutes and complementary) when analysing competing two-sided marketplaces. As we discussed in Section 1.2, strategies of marketplaces and traders will be affected by these factors. Moreover, in fact, currently no work has considered traders' strategies across multiple double auction marketplaces in the multi-home trading environment (although trading strategies across multiple single-sided auction indeed exist, and heuristic bidding strategies used for isolated double auctions can be used in the single-home trading environments with multiple double auction marketplaces). In Chapter 4, we will address this challenge by considering different trading environments and different good properties.

Finally, as we saw in Section 2.3.2.2, a considerable body of work exists on the market policy design of an isolated double auction marketplace. However, we do not know which market policy performs well when competing with other policies. Furthermore, the performance of different market policies are affected by traders' behaviour. Currently, there is no systematic work on analysing the performance of market policies in the competing environments by considering different behaviours for traders. To address this research challenge, in Chapter 5, we experimentally analyse how different market policies influence the performance of competing marketplaces in different environments where different bidding strategies are used. Furthermore, based on insights obtained from this analysis, we design our market policies in the CAT competition.

Chapter 3

Analysis of Competing Marketplaces with Discrete Trader Types

In this chapter, we theoretically analyse the market selection strategies for traders and charging strategies for marketplaces in the context of multiple competing double auction marketplaces. This work addresses the research challenges 1 (analysing market selection strategies for traders) and 4 (analysing charging strategies for marketplaces) in Section 1.2. The rationale for so doing is to provide useful insights to guide the design of a charging strategy. As we discussed in Section 1.3, the strategies of traders and marketplaces are affected by each other, and thus we use game theory to analyse the Nash equilibrium strategies. In this chapter, we focus on the single-home trading environment where traders can only enter one marketplace at a time. Such an environment is highly competitive because the marketplaces have to compete fiercely to fully attract a trader (compared to a multi-home trading environment where traders can participate in multiple marketplaces at a time). For this setting, we are interested in analysing how both traders and marketplaces behave strategically. We will analyse the settings with multi-home trading and hybrid trading in Chapter 4.

The structure of this chapter is as follows. Firstly, we propose a game-theoretic framework for analysing competing marketplaces (Section 3.1). Then, based on the framework, we analyse the Nash equilibrium market selection strategies for traders that are restricted to a predefined number of discrete types (Section 3.2). After having established the traders' equilibrium strategies, we go on to use two different approaches to analyse equilibrium charging strategies for marketplaces (Section 3.3). Finally, we summarise in Section 3.4.

3.1 A Game-Theoretic Framework for Competing Marketplaces

In this section we introduce the game-theoretic framework which models a setting with competing marketplaces and forms the basis of our analysis. Specifically, in our framework, a trading round is considered[1] and it proceeds as follows. First, all marketplaces publish their fees. Second, based on the observed fees, each trader selects a marketplace according to its market selection strategy. Third, traders submit their shouts according to their bidding strategies. Finally, after all traders have submitted their shouts, the marketplace matches buyers and sellers and then executes transactions. In the following, we start by introducing the basic notation of our framework. Then we introduce the marketplaces' charging strategies and the traders' market selection strategies respectively. Finally, we provide the definition of equilibrium strategies in the context of our system.

3.1.1 Buyers and Sellers

We consider a set of buyers, $\mathcal{B} = \{1, 2, ...B\}$, and a set of sellers, $\mathcal{S} = \{1, 2, ...S\}$. Each buyer is interested in purchasing one item, and each seller has one item for sale. All items are identical. Each buyer and seller has a type, which is denoted as θ^b and θ^s respectively. The type of a buyer denotes its *limit price*, i.e. the highest price it is willing to buy the item for, and the type of a seller denotes its *cost price*, i.e. the lowest price it is willing to sell the item for (see Section 2.2). We assume that the types of all buyers are independently drawn from the same cumulative distribution function F^b, with support $[\underline{l}, \bar{l}]$, and the types of all sellers are independently drawn from the same cumulative distribution function F^s, with support $[\underline{c}, \bar{c}]$. 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. However, the type of each specific trader is not known to the other traders or the marketplaces. In addition, we assume that there is a set of competing marketplaces $\mathcal{M} = \{1, 2, ...M\}$, that offer places for trade and provide a centralised matching service between the buyers and sellers.

3.1.2 Marketplaces and Fees

Since we consider marketplaces to be commercial enterprises that seek to make a profit, we assume that they charge fees for their service as match makers. Recall that in Section 1.1, we introduced two types of fees which are common in the real-world marketplaces: *ex ante* fees, which are charged to traders before they make transactions (e.g. the listing fees adopted by eBay), and *ex post* fees, which are charged to traders after they successfully make transactions (e.g. the final fees adopted by eBay and Amazon). In this analysis, we consider registration and profit fees which are typical examples of such fees. In more detail, we define a fee structure of a marketplace m to be the tuple $p_m = (r_m, q_m) \in \mathcal{P}$, $r_m \geq 0$ and $q_m \in [0, 1]$, where r_m is a

[1] This means that we consider a one-shot game. The competition between multiple double auctions is usually regarded as a repeated game. However, the analysis on this game is very complicated, and currently there is no existing work or models we can build upon. Therefore, in this thesis, we look at a one-shot game. In the future work, we would like to extend the analysis to the repeated game. However, we note that the repeated game consists of single one-shot games, and any sequence of one-shot game Nash equilibria is a subgame-perfect equilibrium of the repeated game (Fudenberg and Tirole, 1991). Therefore, insights from analysing the one-shot game will be also useful for designing a competing double auction marketplace.

fixed registration fee charged to a trader when it enters the marketplace, q_m is the percentage fee charged on the trader's profit, which is the difference between the trader's shout and the transaction price, and in the following, we refer to such a fee as a profit fee, and $\mathcal{P}$ is the set of all allowable fee structures. Then the fee structures of all competing marketplaces constitute a fee system $\bar{P} = \langle p_1, p_2, ...p_M \rangle \in \mathcal{P}^M$, where $\mathcal{P}^M$ is the set of all allowable fee systems. Now we describe how a marketplace will set its fee structure. In this work, we consider a mixed charging strategy, where each fee structure is selected with some probability. A pure strategy can be regarded as a degenerate case of a mixed strategy, where the particular pure strategy is selected with probability 1 and every other strategy with probability 0. Now, a mixed charging strategy of marketplace m is defined as $\mu_m : \mathcal{P} \rightarrow [0, 1]$, which means that the probability that marketplace m sets fee structure p_m is $\mu_m(p_m)$, where $\sum_{p_m \in \mathcal{P}} \mu_m(p_m) = 1$. We use $\bar{\mu} = \langle \mu_1(\cdot), ...\mu_M(\cdot) \rangle$ to represent the charging strategy profile of all marketplaces. In addition, we use $\mu_{-m}(\cdot)$ to represent the charging strategy profile of all marketplaces except for marketplace m. Then we can rewrite $\bar{\mu}$ as $\bar{\mu} = \langle \mu_m(\cdot), \mu_{-m}(\cdot) \rangle$.

Finally, we use the k-pricing policy to determine the transaction price of a matched buyer and seller (see Section 2.3.2.2), where the transaction price of a successful interaction in marketplace m is determined by a pricing parameter $k_m \in [0, 1]$, which sets the transaction price of a matched buyer and seller at the point determined by k_m in the interval between their shouts. For example, when a bid d^b is matched with an ask d^s in marketplace m, the transaction price is

$$\text{TP} = k_m * d^s + (1 - k_m) * d^b \tag{3.1}$$

The pricing parameters of all marketplaces constitute the pricing system $\bar{K} = \langle k_1, k_2, ..., k_M \rangle$[2].

3.1.3 Traders' Market Selection Strategies

After describing the charging strategies of the marketplaces, we now introduce the traders' market selection strategies. We assume that each trader has a mixed market selection strategy, whereby each marketplace is selected with some probability. Furthermore, since the fees are determined before the traders choose their marketplaces, the strategy is a function of the fee system. Specifically, a mixed market selection strategy of buyer i is defined as a function $\omega_i^b : [\underline{l}, \bar{l}] \times \mathcal{M} \times \mathcal{P}^M \rightarrow [0, 1]$, where $\omega_i^b(\theta^b, m, \bar{P})$ denotes the probability that buyer i with type $\theta^b \in [\underline{l}, \bar{l}]$ chooses the marketplace $m \in \mathcal{M}$ given the fee system $\bar{P}$, satisfying $\sum_{m \in \mathcal{M}} \omega_i^b(\theta^b, m, \bar{P}) \leq 1$. Here, $1 - \sum_{m \in \mathcal{M}} \omega_i^b(\theta^b, m, \bar{P})$ is the probability that buyer i with type θ^b chooses no marketplace. This happens when buyer i finds it has a negative expected profit in each marketplace. We use $\bar{\omega}^b(\bar{P}) = \langle \omega_1^b(\cdot, \bar{P}), ..., \omega_B^b(\cdot, \bar{P}) \rangle$ to represent the strategy profile of all buyers in the fee system $\bar{P}$. In addition, we use $\omega_{-i}^b(\cdot, \bar{P})$ to represent the strategy profile of all buyers except i. Then $\bar{\omega}^b(\bar{P})$ can be rewritten as $\bar{\omega}^b(\bar{P}) = \langle \omega_i^b(\cdot, \bar{P}), \omega_{-i}^b(\cdot, \bar{P}) \rangle$. Similarly, we use $\omega_j^s : [\underline{c}, \bar{c}] \times \mathcal{M} \times \bar{P} \rightarrow [0, 1]$ to define the probability of selecting a marketplace of seller j and

[2]In this work we use a fixed pricing policy and so it does not form part of the strategy of a marketplace, but the framework can be easily extended to include this as part of the strategy as well.

use $\bar{\omega}^s(\bar{P}) = \langle \omega_1^s(\cdot, \bar{P}), ..., \omega_S^s(\cdot, \bar{P}) \rangle$ to represent the strategy profile of all sellers in the fee system $\bar{P}$, and rewrite it as $\bar{\omega}^s(\bar{P}) = \langle \omega_j^s(\cdot, \bar{P}), \omega_{-j}^s(\cdot, \bar{P}) \rangle$.

3.1.4 Definition of Equilibrium Strategies for Selecting a Marketplace and Setting Fees

Before we can analyse how traders actually select marketplaces and how competing marketplaces set fees (which we discuss in Section 3.2 and 3.3 respectively), we first need to specify the expected utility functions for traders and marketplaces, and define an appropriate solution concept in the context of competing marketplaces.

To this end, we first describe a buyer's expected utility equation for a given fee system $\bar{P}$. A seller's expected utility can be given analogously. Given a buyers' strategy profile $\bar{\omega}^b(\bar{P})$ and a sellers' strategy profile $\bar{\omega}^s(\bar{P})$ in the fee system $\bar{P}$, the expected utility of a buyer i with type θ^b in the fee system $\bar{P}$ and pricing system $\bar{K}$ is defined by:

$$\tilde{U}_i^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) = \sum_{m=1}^{M} \omega_i^b(\theta^b, m, \bar{P}) \times \tilde{U}_{i,m}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) \tag{3.2}$$

where $\tilde{U}_{i,m}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b)$ is buyer i's expected utility if it chooses to trade in marketplace m.

Furthermore, marketplace m's expected utility given a charging strategy profile $\bar{\mu}$, pricing parameter k_m and traders' market selection strategy profiles $\bar{\omega}^b(\cdot)$ and $\bar{\omega}^s(\cdot)$, is as follows:

$$\tilde{U}_m(\bar{\mu}) = \sum_{\bar{P} \in \mathcal{P}^M} \mu(\bar{P}) * \tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P})) \tag{3.3}$$

where $\mu(\bar{P}) = \prod_{m \in \mathcal{M}} \mu_m(p_m)$ is the probability that fee system $\bar{P} = \langle p_1, ..., p_M \rangle$ is selected, and $\tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}))$ is marketplace m's expected utility given fee system $\bar{P}$. Note that both the buyer's expected utility in marketplace m, $\tilde{U}_{i,m}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b)$, and marketplace m's expected utility, $\tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}))$, depend on the specific bidding strategies adopted by traders and the matching policy adopted by marketplace m. We will detail them in Section 3.2.1 and 3.3.2.2 respectively where we consider a particular market setting.

After providing general forms for the traders and the marketplaces' expected utilities, we are now ready to define the equilibrium strategies for traders and marketplaces in our system. Since we consider a game with incomplete information about traders' types, the Bayes-Nash equilibrium (BNE) solution concept (see Section 2.1), in which each player's strategy maximises its expected utility given other players' strategies, is the most appropriate to define this equilibrium behaviour. Here, we define equilibrium strategies of both traders and marketplaces as a whole, since in our system, traders' market selection strategies and marketplaces' charging strategies affect each other. Formally, the mixed Bayes-Nash equilibrium in our setting is defined as:

Definition *Given pricing system $\bar{K}$, a charging strategy profile $\bar{\mu}^*$ and market selection strategy profiles $\bar{\omega}^{b*}(\cdot)$ and $\bar{\omega}^{s*}(\cdot)$ constitute a mixed Bayes-Nash equilibrium, if:*

$$\forall i \in \mathcal{B}, \forall \theta^b \in [\underline{l}, \bar{l}], \forall \bar{P} \in \mathcal{P}^M, \forall \omega_i^b(\cdot, \bar{P}) \in \Delta^T :$$

$$\tilde{U}_i^b(\bar{P}, \bar{K}, \langle \omega_i^{b*}(\cdot, \bar{P}), \omega_{-i}^{b*}(\cdot, \bar{P}) \rangle, \bar{\omega}^{s*}(\bar{P}), \theta^b) \geq \tilde{U}_i^b(\bar{P}, \bar{K}, \langle \omega_i^b(\cdot, \bar{P}), \omega_{-i}^{b*}(\cdot, \bar{P}) \rangle, \bar{\omega}^{s*}(\bar{P}), \theta^b);$$

i.e. each buyer's strategy is a best response to other traders' strategies for each possible fee system.

and $\forall j \in \mathcal{S}, \forall \theta^s \in [\underline{c}, \bar{c}], \forall \bar{P} \in \mathcal{P}^M, \forall \omega_j^s(\cdot, \bar{P}) \in \Delta^T :$

$$\tilde{U}_j^s(\bar{P}, \bar{K}, \bar{\omega}^{b*}(\bar{P}), \langle \omega_j^{s*}(\cdot, \bar{P}), \omega_{-j}^{s*}(\cdot, \bar{P}) \rangle, \theta^s) \geq \tilde{U}_j^s(\bar{P}, \bar{K}, \bar{\omega}^{b*}(\bar{P}), \langle \omega_j^s(\cdot, \bar{P}), \omega_{-j}^{s*}(\cdot, \bar{P}) \rangle, \theta^s);$$

i.e. each seller's strategy is a best response to other traders' strategies for each possible fee system.

and $\forall m \in \mathcal{M}, \forall \mu_m(\cdot) \in \Delta^M :$

$$\tilde{U}_m(\langle \mu_m^*(\cdot), \mu_{-m}^*(\cdot) \rangle) \geq \tilde{U}_m(\langle \mu_m(\cdot), \mu_{-m}^*(\cdot) \rangle)$$

i.e. each marketplace's charging strategy is a best response to other marketplaces' charging strategies.

where Δ^T is the set of all possible (mixed) market selection strategies and Δ^M is the set of all possible (mixed) charging strategies.

Given the equilibrium definition, in what follows, we will analyse the both the traders' equilibrium market selection strategies (in Section 3.2) and the marketplaces' equilibrium charging strategies (in Section 3.3).

3.2 Equilibrium Analysis of the Market Selection Strategies

After describing the framework for analysing competing double auction marketplaces, we now use this to analyse the equilibrium strategies of market selection for the trading agents. Before doing this, however, we first need to specify the bidding strategies adopted by the traders and the matching policies adopted by the marketplaces. In this analysis, we make a simplifying assumption that traders use a truth-telling bidding strategy, which means that they will submit their types as their shouts during the trading process[3]. For the matching policy, we consider the *equilibrium matching* policy (see Section 2.3.2.2) since this aims to maximise traders' profits and thus maximises the allocative efficiency for the marketplace. Given the specific bidding strategy and matching policy, in the following, we will derive traders' expected utilities in this setting, and then game-theoretically and dynamically analyse traders' equilibrium strategies of market selection for a given fee system. We are interested in calculating the *symmetric* Bayes-Nash equilibria (BNEs), as is common in game theory for settings with incomplete information,

[3]As we said in Section 1.2, deriving the traders' equilibrium bidding strategies in double auctions is a challenging problem. Furthermore, in this chapter, we want to focus the relationship between marketplaces' fees and traders' market selection. Therefore, here we consider a simple bidding strategy, which can be easily mathematically represented, to simplify the analysis. While the bidding strategies can affect traders' expected profits, and in turn affect their market selection, the analysis based on this bidding strategy will still provide us valuable insights about the relationship between marketplaces' fees and traders' market selection strategies. In Chapter 4, we will analyse traders' equilibrium bidding strategies by using fictitious play.

and so we can assume that (in equilibrium) traders with the same type will employ the same strategy. Thus in the following equations, we omit the indexes i and j when referring to specific buyers and sellers, and we add the convention that the market selection strategy profile consists of the market selection strategy of each type, instead of each trader.

3.2.1 A Trader's Expected Utility

In what follows, we derive the expected utility of a buyer with type θ^b in the fee system $\bar{P}$ given the market selection strategy profiles of buyers and sellers, $\bar{\omega}^b(\bar{P})$ and $\bar{\omega}^s(\bar{P})$. The seller's expected utility is calculated analogously. According to Equation 3.2, we need to calculate the trader's expected utility in each marketplace m. Intuitively, the trader's expected utility in marketplace m not only depends on its own type, but also on the number and types of other traders choosing this marketplace. While F^b and F^s are the overall type distribution functions of buyers and sellers respectively, the distribution of types within an individual marketplace can differ depending on which types choose what marketplace. We refer to the type distribution of a specific marketplace m as the local type distribution, which is also a cumulative distribution function. This is derived as follows. For a given fee system $\bar{P}$ and market selection strategy $\omega^b(\cdot)$, the probability that the type of a buyer is less than θ^b in marketplace m is:

$$H^b_m(\theta^b|\bar{P}) = \int_{\underline{l}}^{\theta^b} f^b(x) * \omega^b(x, m, \bar{P})dx \tag{3.4}$$

We can see that the prior probability that a buyer (irrespective of its type) will choose marketplace m is given by $H^b_m(\bar{l}|\bar{P})$. Note that the above function is not a proper local type distribution function because the function range may be smaller than 1. Then, to obtain a proper local type distribution function of buyers in marketplace m, we need to normalise the above equation:

$$G^b_m(\theta^b|\bar{P}) = \frac{H^b_m(\theta^b|\bar{P})}{H^b_m(\bar{l}|\bar{P})} \tag{3.5}$$

This is also called the local type distribution. Furthermore, the local probability density function of buyer types is:

$$g^b_m(\theta^b|\bar{P}) = \frac{f^b(\theta^b) * \omega^b(x, m, \bar{P})}{H^b_m(\bar{l}|\bar{P})} \tag{3.6}$$

The equations of the sellers can be derived in the same way.

Recall that, in addition to the types of traders, the expected utility also depends on the (expected) number of traders choosing this marketplace. To this end, we calculate the probabilities that there are *exactly* τ^b other buyers (excluding the buyer for which we are calculating the expected utility) and τ^s sellers choosing marketplace m, which are given by the binomial distributions:

$$\rho^b_m(\tau^b) = \binom{B-1}{\tau^b} * \left(H^b_m(\bar{l}|\bar{P})\right)^{\tau^b} * \left(1 - H^b_m(\bar{l}|\bar{P})\right)^{B-1-\tau^b} \tag{3.7}$$

$$\rho_m^s(\tau^s) = \binom{S}{\tau^s} * \left(H_m^s(\bar{c}|\bar{P})\right)^{\tau^s} * \left(1 - H_m^s(\bar{c}|\bar{P})\right)^{S-\tau^s} \tag{3.8}$$

We now proceed to calculate the buyer's expected utility, which is affected by the matching policy of marketplaces. To this end, recall that in the *equilibrium matching* policy, the marketplace matches the buyer with the v-th highest limit price with the seller with the v-th lowest cost price. We then calculate the probability that the buyer's type θ^b is at a certain position. Specifically, for given fee system $\bar{P}$, when $\tau^b + 1$ buyers choose marketplace m, the probability that the buyer with type θ^b is the v-th ($v = 1, ..., \tau^b + 1$) highest is given by:

$$Pr_m^b(v|\theta^b, \bar{P}) = \binom{\tau^b}{v-1} * \left(1 - G_m^b(\theta^b|\bar{P})\right)^{v-1} * \left(G_m^b(\theta^b|\bar{P})\right)^{\tau^b+1-v} \tag{3.9}$$

Similarly, the probability that the seller's type θ^s is the v-th ($v = 1, ..., \tau^s$) lowest among τ^s sellers in marketplace m is given by:

$$Pr_m^s(v|\theta^s, \bar{P}) = \binom{\tau^s - 1}{v-1} * \left(G_m^s(\theta^s|\bar{P})\right)^{v-1} * \left(1 - G_m^s(\theta^s|\bar{P})\right)^{\tau^s-v} \tag{3.10}$$

Furthermore, the prior probability that a seller is the v-th lowest is given by:

$$Pr^s(v|\bar{P}) = \int_{\underline{c}}^{\bar{c}} Pr_m^s(v|\theta^s, \bar{P}) * g_m^s(\theta^s|\bar{P}) d\theta^s \tag{3.11}$$

Now using Bayes' theorem, we can calculate the probability density function of a seller at position v:

$$g_m^s(\theta^s|v, \bar{P}) = \frac{Pr_m^s(v|\theta^s, \bar{P}) * g_m^s(\theta^s|\bar{P})}{Pr^s(v|\bar{P})} \tag{3.12}$$

At this moment, we can get the buyer's expected gross profit (without taking into account any fees that the buyer pays to the marketplace) in marketplace m with strategy profiles of buyers and sellers, $\bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P})$:

$$\tilde{\Lambda}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) = \sum_{\tau^b=0}^{B-1} \rho_m^b(\tau^b) * \sum_{v=1}^{\tau^b+1} Pr_m^b(v|\theta^b, \bar{P}) * \Big(\sum_{\tau^s=v}^{S} \rho_m^s(\tau^s) * \int_{\theta^s=\underline{c}}^{\theta^b} k_m * (\theta^b - \theta^s) * g_m^s(\theta^s|v, \bar{P}) d\theta^s\Big) \tag{3.13}$$

where $\theta^b - \theta^s$ is called the trading surplus, and $k_m * (\theta^b - \theta^s)$ is the share of the buyer's surplus, which is determined by the pricing parameter k_m. By also considering the registration fee and profit fee from marketplace m, a buyer's expected utility in this marketplace becomes:

$$\tilde{U}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) = \tilde{\Lambda}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) * (1 - q_m) - r_m \tag{3.14}$$

The above equation gives the expected utility of the buyer in a particular marketplace. Therefore, a buyer's expected utility over all marketplaces in the fee system $\bar{P}$ is:

$$\tilde{U}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) = \sum_{m=1}^{M} \omega^b(\theta^b, m, \bar{P}) * \tilde{U}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), \theta^b) \tag{3.15}$$

After deriving traders' expected utilities, in the next section, we will game theoretically analyse the equilibrium market selection strategies for traders.

3.2.2 A Game-Theoretic Analysis of Market Selection Strategies

In this section, we analytically derive the traders' equilibrium market selection strategies for a given fee system $\bar{P} = \langle p_1, ..., p_M \rangle$ (i.e. each marketplace $m \in \mathcal{M}$ sets the fee structure p_m with 100% probability ($\mu_m(p_m) = 1$)). As we said before, we focus on the symmetric BNE which means that traders with the same type will adopt the same strategy in equilibrium. Furthermore, in order to get insights from this complicated game with more traders and more types while allowing for tractable results, we initially make several simplifying assumptions (specified below). In the next section, we will use evolutionary game theory to computationally determine the equilibrium which will allow us to relax some of these assumptions.

Specifically, we make the following assumptions. First of all, we consider the competition between two marketplaces, i.e. $M = 2$ (this is consistent with the previous theoretical work introduced in Section 2.4. However, in Section 3.2.3.4, we discuss the setting with more than two competing marketplaces), and we restrict our analysis to two buyers and two sellers, i.e. $B = S = 2$, (although we will relax this in Subsection 3.2.3 and Section 3.3). In addition, we restrict our analysis to discrete trader types. In particular, we assume that there are two types of buyers and two types of sellers: rich and poor, which are denoted by t_2^b and t_1^b respectively for buyers, and t_1^s and t_2^s for sellers. A rich buyer is defined as having a higher limit price than a poor buyer, i.e. $t_2^b > t_1^b$, and a rich seller is defined as having a lower cost price than a poor seller, i.e. $t_1^s < t_2^s$. Trader types are independently drawn from the discrete uniform distribution (i.e. both types are equally likely). In Chapter 4, we will extend this analysis by considering continuous trader types. Furthermore, we only consider profit fees at this stage (i.e. $r_1 = r_2 = 0$). This simplifies the analysis since traders always have non-negative profits when participating, and they will always choose one of the marketplaces. In this way we can reduce the strategy space since $\omega^b(\theta^b, 1, \bar{P}) = 1 - \omega^b(\theta^b, 2, \bar{P})$, and similar for sellers. We will extend our analysis to registration fees in Subsection 3.2.3 and Section 3.3.

Given these assumptions, we now investigate the traders' market selection equilibrium behaviour. Intuitively, we can see that all traders selecting one marketplace constitutes a pure strategy BNE, since given all other traders selecting one marketplace, the best response of a trader is also to select this marketplace (otherwise they will have nobody to trade with). In addition to the pure strategy BNEs, we are also interested in the mixed symmetric BNE of the traders' market selection strategies since we would like to know whether two competing marketplaces can co-exist. Previously, note that we derived a trader's expected utility considering continuous trader types (see Equation 3.15). Now we need to adapt this equation to discrete trader types. However, since this is straightforward, we will not show it in detail. As we know, in the mixed Nash equilibrium, a player should be indifferent between choosing each of the pure strategies that form part of the mixed strategy, i.e. its expected utility for each of these pure strategies should be the same (Osborne and Rubinstein, 1994). Since in our setting each trader has only two pure strategies, i.e. choosing marketplace 1 or 2, in any non-pure Nash equilibrium the traders should be indifferent between these choices. Thus we get the following equations to calculate the mixed BNE (one for each type of buyers and sellers):

$$\tilde{U}_1^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b) = \tilde{U}_2^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b) \quad (3.16)$$

$$\tilde{U}_1^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b) = \tilde{U}_2^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b) \quad (3.17)$$

$$\tilde{U}_1^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^s) = \tilde{U}_2^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^s) \quad (3.18)$$

$$\tilde{U}_1^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^s) = \tilde{U}_2^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^s) \quad (3.19)$$

If the solution is in the range [0, 1], then this constitutes a mixed symmetric BNE.

By expanding the above equations (see Appendix A), we find that the solution such that

$$\omega^b(t_1^b, 1, \bar{P}) = \omega^b(t_2^b, 1, \bar{P}) \quad (3.20)$$

and

$$\omega^s(t_1^s, 1, \bar{P}) = \omega^s(t_2^s, 1, \bar{P}) \quad (3.21)$$

always exists. This means that there always exists a mixed BNE whereby buyers adopt the same mixed strategy no matter whether they are rich or poor, and the same for sellers. In the following, we analyse this solution in more detail to better understand the traders' equilibrium behaviour.

To this end, we first rewrite Equations 3.16-3.19 as the following two equations:

$$\left[2 * \omega^s(t_1^s, 1, \bar{P}) - \left(\omega^s(t_1^s, 1, \bar{P})\right)^2 - \omega^b(t_1^b, 1, \bar{P}) * \omega^s(t_1^s, 1, \bar{P}) + \omega^b(t_1^b, 1, \bar{P}) * \left(\omega^s(t_1^s, 1, \bar{P})\right)^2\right] * k_1 * (1 - q_1)$$
$$= \left[1 - \omega^s(t_1^s, 1, \bar{P}) + \omega^b(t_1^b, 1, \bar{P}) * \omega^s(t_1^s, 1, \bar{P}) - \omega^b(t_1^b, 1, \bar{P}) * \left(\omega^s(t_1^s, 1, \bar{P})\right)^2\right] * k_2 * (1 - q_2) \quad (3.22)$$

$$\left[2 * \omega^b(t_1^b, 1, \bar{P}) - \left(\omega^b(t_1^b, 1, \bar{P})\right)^2 - \omega^b(t_1^b, 1, \bar{P}) * \omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_1^s, 1, \bar{P}) * \left(\omega^b(t_1^b, 1, \bar{P})\right)^2\right] * (1 - k_1) * (1 - q_1)$$
$$= \left[1 - \omega^b(t_1^b, 1, \bar{P}) + \omega^b(t_1^b, 1, \bar{P}) * \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_1^s, 1, \bar{P}) * \left(\omega^b(t_1^b, 1, \bar{P})\right)^2\right] * (1 - k_2) * (1 - q_2) \quad (3.23)$$

Note that the solution for the above two equations depends on how both marketplaces set pricing parameters k_1 and k_2, and how they charge profit fees q_1 and q_2. Now we analyse how these two factors affect the mixed BNEs respectively. Firstly, we assume that both marketplaces charge the same profit fee (from Equations 3.22 and 3.23, we can see that when $q_1 = q_2$, they can be cancelled from both left- and right-hand sides of equations), but have different pricing parameters. Then the resulting mixed BNEs are shown in Figure 3.1(a), from which we find that when the marketplace sets a high value for the pricing parameter, i.e. allocates more profits to buyers and thus less to sellers, buyers have a higher probability of choosing this marketplace, and sellers have a lower probability of choosing it. In this situation, buyers(sellers) have the same expected utility in marketplaces 1 and 2, and thus a mixed BNE is constituted. Then we consider the case where both marketplaces have the same pricing parameter (when $k_1 = k_2$, they

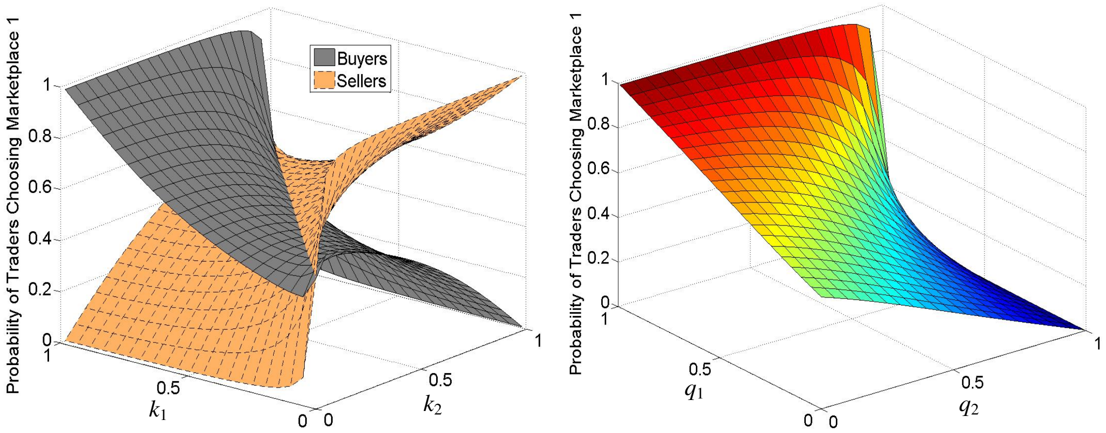

(a) The mixed BNE strategy of traders when both marketplaces charge the same profit fee, but have different pricing parameters. k_1 is the pricing parameter of marketplace 1 and k_2 is the pricing parameter of marketplace 2.

(b) The mixed BNE strategy of traders when both marketplaces have the same pricing parameter, but charge different profit fees. q_1 is the profit fee of marketplace 1 and q_2 is the profit fee of marketplace 2.

FIGURE 3.1: The mixed BNE strategy of traders.

can be cancelled in equations), but charge different profit fees. In this situation, by analysing Equations 3.22 and 3.23, we can see that all traders have the same market selection strategy. The results are shown in Figure 3.1(b). Counter-intuitively, we find that, in equilibrium, traders have a higher probability of choosing marketplace 1 when it charges a higher profit fee. This is because when buyers have beliefs that sellers have a higher probability of choosing marketplace 1, then they will prefer to choose marketplace 1 even though it charges a higher fee, and they still have the same expected utility in marketplaces 1 and 2. Mutatis mutandis for sellers. This implies that it is possible for the competing marketplace to charge a higher fee to make more profits while still maintaining market share at a good level. In the following section, we will analyse this phenomenon in more detail.

3.2.3 An Evolutionary Analysis of Market Selection Strategies

In the above, we game-theoretically analysed the traders' equilibrium behaviour with regards to market selection strategies and showed that there exist at least three BNEs: all traders choosing marketplace 1 or 2 and the mixed BNE. As we discussed in Section 2.1.1, such equilibria only provide a static explanation for why populations playing BNE strategies remain in that state since each population makes a best response to the other populations' strategies. Therefore, this solution concept fails to indicate whether the BNE can be reached and which of these equilibria is most likely to be converged to. To overcome this and to analyse settings with more than 2 buyers and 2 sellers, in what follows, we use evolutionary game theory (EGT) to analyse this game, which focuses on the dynamic change of strategies rather than the static properties of Nash equilibria (see Section 2.1.1). In EGT, players gradually adjust their strategies over time in response to the repeated observation of their opponents' strategies[4]. In the following, we first

[4]Note that, although this process is a repeated learning process, the game itself is not a repeated game, but a one-shot game. We just use such a repeated learning process to analyse how traders learn to change their strategies,

describe the replicator dynamics equations used in our analysis, which capture the dynamics of traders' market selection strategies, and then give the evolutionary analysis in detail.

3.2.3.1 Replicator Dynamics

In EGT, the *replicator dynamics* equation is often used to specify the dynamic adjustment of the probability of which pure strategy should be played (see Section 2.1.1). In our work, we consider 4 different types of traders (rich buyers, poor buyers, rich sellers and poor sellers). In order to allow different types of traders to converge to different equilibrium strategies, we use a different population to evolve the strategy of each type. Therefore first we introduce the 4-population replicator dynamics equations which show the dynamic changes of traders' market selection strategies with respect to time t:

$$\begin{aligned}
\dot{\omega}^b(t_1^b, m, \bar{P}) &= \frac{d\omega^b(t_1^b, m, \bar{P})}{dt} \\
&= (\tilde{U}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b) - \tilde{U}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b)) * \omega^b(t_1^b, m, \bar{P}) \quad (3.22) \\
\dot{\omega}^b(t_2^b, m, \bar{P}) &= \frac{d\omega^b(t_2^b, m, \bar{P})}{dt} \\
&= (\tilde{U}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b) - \tilde{U}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b)) * \omega^b(t_2^b, m, \bar{P}) \quad (3.23) \\
\dot{\omega}^s(t_1^s, m, \bar{P}) &= \frac{d\omega^s(t_1^s, m, \bar{P})}{dt} \\
&= (\tilde{U}_m^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^s) - \tilde{U}^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^s)) * \omega^s(t_1^s, m, \bar{P}) \quad (3.24) \\
\dot{\omega}^s(t_2^s, m, \bar{P}) &= \frac{d\omega^s(t_2^s, m, \bar{P})}{dt} \\
&= (\tilde{U}_m^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^s) - \tilde{U}^s(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^s)) * \omega^s(t_2^s, m, \bar{P}) \quad (3.25)
\end{aligned}$$

Note that the 4 populations interact through the utility functions, which depend on the strategies of other trader types. As an example, $\dot{\omega}^b(t_1^b, m, \bar{P})$ describes how the poor buyer with type t_1^b changes its probability of choosing marketplace m in the fee system $\bar{P}$. Here, $\tilde{U}_m^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b)$ is the poor buyer's expected utility when choosing marketplace m given market selection strategy profiles $\bar{\omega}^b(\bar{P})$ and $\bar{\omega}^s(\bar{P})$, and $\tilde{U}^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_1^b)$ is the poor buyer's overall expected utility (see Subsection 3.2.1). In order to get the dynamics of the strategies, we need to calculate *trajectories*, which indicate how the mixed strategies evolve. In more detail, initially, a mixed strategy is chosen as a starting point (in our results, we experiment with a large number of points). In the following, we denote a starting point by $(\omega^b(t_2^b, 1, \bar{P}), \omega^b(t_1^b, 1, \bar{P}), \omega^s(t_1^s, 1, \bar{P}), \omega^s(t_2^s, 1, \bar{P}))$. The dynamics are then calculated according to the above replicator equations. According to the dynamic changes of traders' strategies, their current mixed strategies can be calculated. Such calculations are repeated until $\dot{\omega}^b(\cdot)$ and $\dot{\omega}^s(\cdot)$ become zero, at which point the equilibrium is reached. When considering traders evolving from all possible starting points, we get several regions. The region from which all trajectories converge to a particular equilibrium is called the *basin of attraction* of this equilibrium. The basin is very useful since given the assumption that each starting point is selected by traders with an equal probability, its size can be used as an indicator of the probability of traders converging to

and thus, which equilibrium, if any, can be reached.

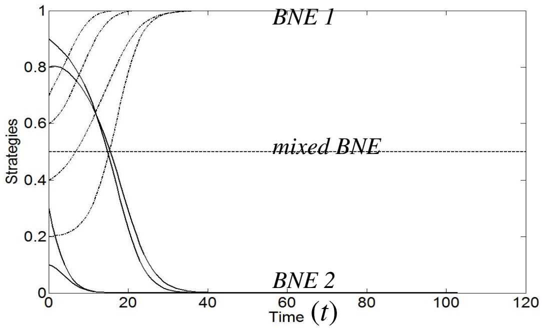

(a) Equilibrium behaviour of 2 buyers and 2 sellers with three chosen starting points.

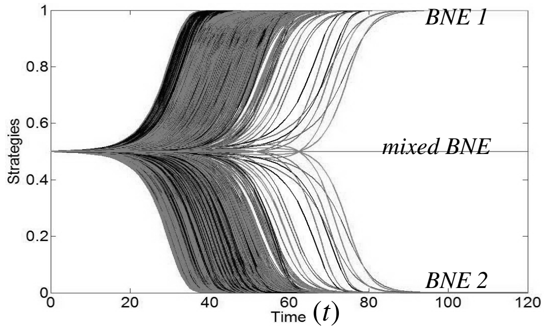

(b) Equilibrium behaviour of 2 buyers and 2 sellers with starting points around (0.5, 0.5, 0.5 0.5).

FIGURE 3.2: Equilibrium behaviour of 2 buyers and 2 sellers in the setting with 2 identical marketplaces when $q_1 = q_2 = 10\%$ and $r_1 = r_2 = 0$.

that equilibrium. However, we should note that replicator dynamics equations are only used to find which equilibrium traders are likely to converge to, and they do not show the realistic way that traders select marketplaces.

3.2.3.2 Experimental Results

After providing replicator dynamics equations, we now analyse how traders dynamically evolve their market selection strategies to converge to the equilibrium.

4-population with two identical marketplaces:

Firstly, we analyse the general cases with four different populations (rich buyers, poor buyers, rich sellers and poor sellers). For illustrative purposes, we assign the traders' types as follows: $t_1^b = 4$, $t_2^b = 8$, $t_1^s = 0$ and $t_2^s = 3$.[5] At this stage, we assume that both marketplaces only charge profit fees. We first consider the case where the two competing marketplaces are identical. Specifically, we consider $k_1 = k_2 = 0.5$ and $q_1 = q_2 = 10\%$ as an example. According to Equations 3.22 and 3.23, we know that (0.5, 0.5, 0.5, 0.5) is a mixed BNE. We now show the evolutionary results of two representative starting points (0.6, 0.4, 0.7, 0.2), (0.1, 0.8, 0.3, 0.9) and a specific starting point (0.5, 0.5, 0.5, 0.5) in Figure 3.2(a). The x-axis is the time at which the mixed strategies evolve, the points at $t = 0$ correspond to the starting points, from which traders evolve their strategies. As can be seen, traders eventually converge to BNE 1 (i.e. marketplace 1) or BNE 2 (i.e. marketplace 2) except that when starting at the mixed BNE, traders will stay at the mixed BNE. Furthermore, we analyse the area of the starting points around the equilibrium (0.5, 0.5, 0.5, 0.5) and find that these do not converge to (0.5, 0.5, 0.5, 0.5). This is described in Figure 3.2(b), which shows the evolutionary results when starting points are chosen from 0.499 to 0.501 with step size 0.0005. Therefore, we can conclude that even though (0.5, 0.5, 0.5, 0.5) is a mixed BNE, it is a saddle point where no trajectories converge, and is unlikely to be reached.

[5]Other type values can be chosen. However, our experiment analysis confirms that the conclusions are similar.

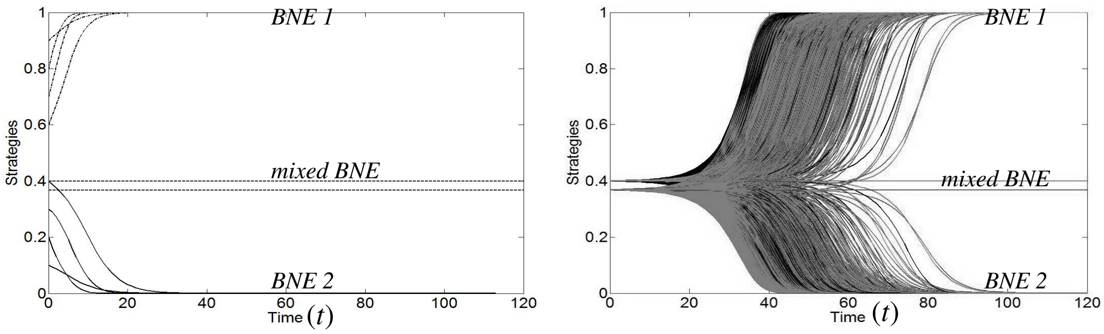

(a) Equilibrium behaviour of 2 buyers and 2 sellers with three chosen starting points.

(b) Equilibrium behaviour of 2 buyers and 2 sellers with starting points around (0.3679, 0.3679, 0.3994, 0.3994).

FIGURE 3.3: Equilibrium behaviour of 2 buyers and 2 sellers in the setting with 2 different marketplaces when $q_1 = 20\%$, $q_2 = 40\%$ and $r_1 = r_2 = 0$.

4-population with two different marketplaces:

Now we consider the traders' evolved strategies when fees and pricing parameters are different across marketplaces. As an example, we let $k_1 = 0.48$, $k_2 = 0.51$, $q_1 = 20\%$ and $q_2 = 40\%$. By solving Equations 3.22 and 3.23, we find one mixed BNE at (0.3679, 0.3679, 0.3994, 0.3994). Then we show the evolutionary results of two representative starting points (0.2, 0.4, 0.3, 0.1), (0.7, 0.9, 0.8, 0.6), and the specific starting point (0.3679, 0.3679, 0.3994, 0.3994) in Figure 3.3(a) and the evolutionary results with starting points around the equilibrium (0.3679, 0.3679, 0.3994, 0.3994) in Figure 3.3(b). We still find that traders finally converge to either marketplace 1 or 2, and the mixed BNE is unlikely to occur. Furthermore, we also ran experiments from a large number of starting points, and still find that traders eventually converge to one marketplace in equilibrium.

2-population:

In the above, we have analysed the dynamics of traders' market selection strategies with 4 populations. However, in these cases, it is difficult to *clearly* visualise how traders evolve to converge to the equilibrium when considering *a variety of starting points*. Therefore, in order to be able to further illustrate the dynamics of EGT when traders evolve from various starting points, we run all next sets of experiments by assuming that rich buyers and rich sellers have the same behaviour, and poor buyers and poor sellers have the same behaviour. In order to reasonably make this assumption, rich(poor) buyers and rich(poor) sellers should be treated equally by marketplaces. Thus, we assume pricing parameter $k_m = 0.5$, i.e. the transaction price is set in the middle of shouts of the matched buyers and sellers, which means that the marketplaces have no bias in favor of buyers or sellers when allocating surpluses. Furthermore, we assume that surpluses of buyers and sellers are symmetric and, as an example, we let $t_1^b = 4$, $t_2^b = 6$, $t_1^s = 0$ and $t_2^s = 2$. Based on these assumptions, in the following experiments, Equations 3.22 and 3.25 are equivalent, as are Equations 3.23 and 3.24. By so doing, we reduce the 4-population replicator dynamics to 2-population, and can visualise traders' dynamics of market selection strategies in 2-dimensional graphs.

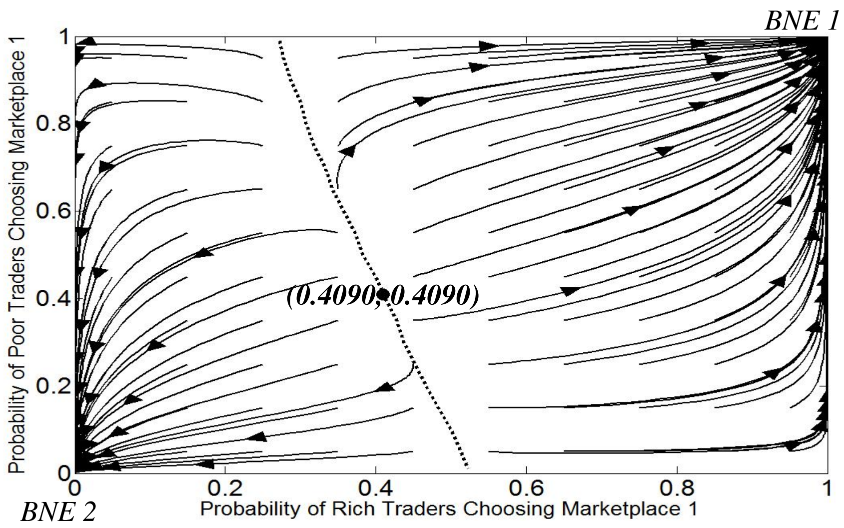

FIGURE 3.4: Evolutionary process of market selection strategies with 2 buyers and 2 sellers when q_1 = 50%, q_2 = 60% and $r_1 = r_2 = 0$. The dotted line denotes the boundary between the basins of attractions.

2-population with 2 buyers and 2 sellers:

Firstly, we still consider 2 buyers and 2 sellers. We assume that marketplace 1 charges 50% profit fee and marketplace 2 charges 60% profit fee. Then the mixed BNE satisfying Equations 3.20 and 3.21 is (0.4090, 0.4090). The evolutionary results are shown in Figure 3.4, where the x-axis is the rich buyer(seller)'s probability of choosing marketplace 1, and the y-axis is the poor buyer(seller)'s probability of choosing marketplace 1. We find that all traders either converge to BNE 1 (i.e. marketplace 1) or BNE 2 (i.e. marketplace 2) in equilibrium depending on the initial starting points, and no trajectory converges to (0.4090, 0.4090) (the solid circle in Figure 3.4), i.e. the mixed BNE is a saddle point. This indicates that the mixed BNE is hard to reach. Furthermore, the figure also shows that the basin of attraction to BNE 1 is bigger, which means that traders have a higher probability of converging to marketplace 1 since this marketplace charges less than marketplace 2. From the results, we also find that two competing marketplaces cannot co-exist in equilibrium. We then try various fee systems with different combinations of profits fees charged by marketplaces, and still find that all traders converge to one marketplace in equilibrium.

2-population with 5 buyers and 5 sellers:

Now, we extend the above analysis to the case with 5 buyers and 5 sellers (in the following analysis, unless mentioned otherwise, we always assume that there are 5 buyers and 5 sellers). The same as above, we still assume that both marketplaces only charge profit fees. For example, we assume that marketplace 1 charges 20% profit fee and marketplace 2 charges 30% profit fee. The dynamic results for different starting points are shown in Figure 3.5. We still find that all traders eventually converge to one marketplace. By experimenting with various fee systems with different combinations of profit fees, we still find that all traders converge to one marketplace in equilibrium. We then extend our analysis to the case that both competing marketplaces only charge registration fees, and still find that traders eventually converge to one marketplace in

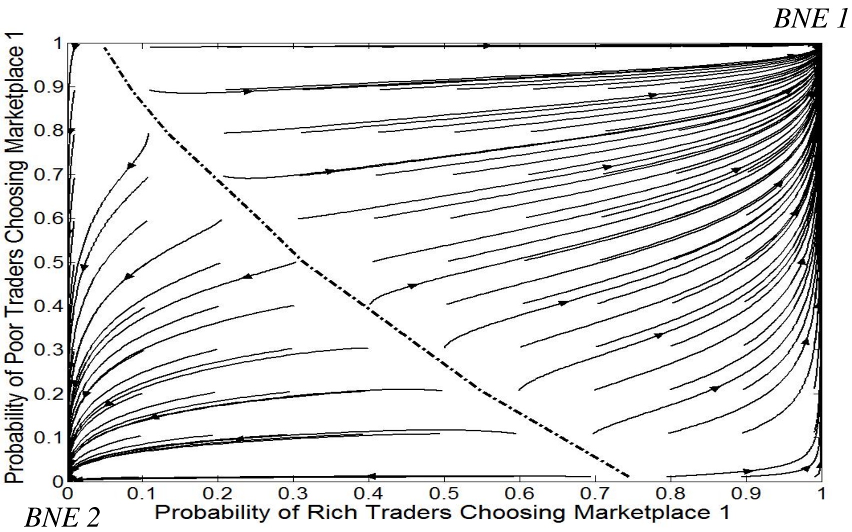

FIGURE 3.5: Evolutionary process of market selection strategies with 5 buyers and 5 sellers when $q_1 = 20\%$, $q_2 = 30\%$ and $r_1 = r_2 = 0$. The dotted line denotes the boundary between the basins of attractions

equilibrium if they want to select a marketplace[6]. These show that in our framework, the positive size effect has a larger impact than the negative size effect, which will cause traders to converge to one marketplace in equilibrium.

2-population with different marketplaces charging different types of fees:

Then we consider the case that different types of fees are charged by competing marketplaces (i.e. marketplace 1 charges profit fees, and marketplace 2 charges registration fees). In this case, we find that for some fee systems (when the profit fee of marketplace 1 is higher than 10%, and the registration fee of marketplace 2 is higher than 0.4), traders may converge to different marketplaces in equilibrium when evolving from certain starting points. As an example, we assume that marketplace 1 charges 50% profit fee and marketplace 2 charges 0.8 registration fee. The results are shown in Figure 3.6. We found that, in this case, when traders evolve from certain starting points, they will converge to BNE 3, where rich traders converge to marketplace 1 which charges a registration fee, and poor traders converge to marketplace 2 which charges a profit fee. At this moment, two competing marketplaces co-exist. In contrast to Ellison et al. (2004), where co-existence of competing marketplaces is caused by negative size effect, here the co-existence is caused by the strong differentiation of competing marketplaces by setting different types of fees. In more detail, rich traders prefer the marketplace charging a lump sum fee since this fee is likely to be smaller than the absolute extracted profit obtained from charging profit fees on a large transaction profit. However, poor traders prefer the marketplace which charges a profit fee, since this can guarantee non-negative profits for them, and a high registration fee may lead to negative profits.

[6]Note that when high registration fees are charged, eventually, poor traders (and rich traders) may not choose any marketplace since high registration fees may cause negative profits for them.

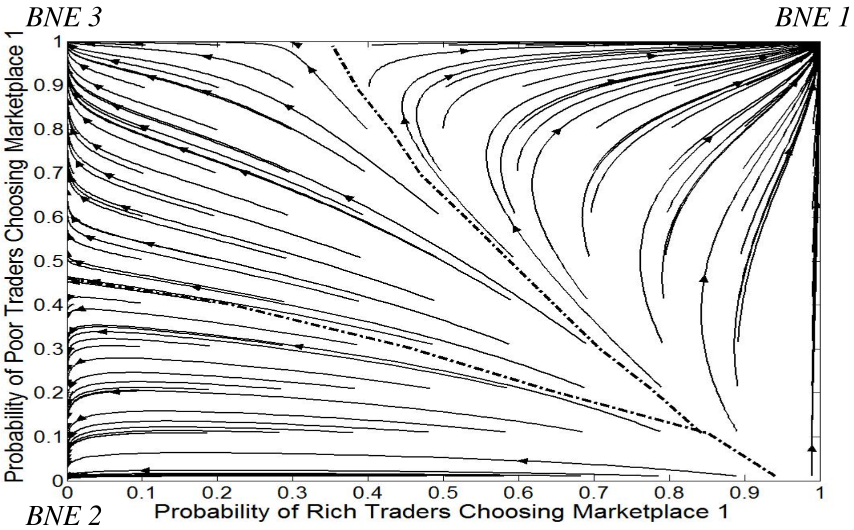

FIGURE 3.6: Evolutionary process of market selection strategies with 5 buyers and 5 sellers when $q_1 = 50\%$, $q_2 = 0$, $r_1 = 0$ and $r_2 = 0.8$. The dotted line denotes the boundary between the basins of attractions.

3.2.3.3 Lock-in Region

In the above analysis where both marketplaces charge profit fees, we find that when evolving from certain starting points, traders may converge to BNE2 (i.e. marketplace 2, which is the more expensive one). This is interesting since it means that the marketplace can charge higher fees to make more profits but still keep traders (even if the size of the basin of attraction is smaller when fees are relatively higher). In the following, we analyse this phenomenon in detail. In doing so, we consider profit fees as an example (i.e. both marketplaces charge no registration fees). Since EGT only works with a finite set of strategies, we discretize the profit fees of the marketplaces. Specifically, we discretize the continuous profit fee from 0 and 1 with a step size of 0.1 (it can also be discretized with other step sizes, such as 0.05). Furthermore, in the following analysis, we still assume that there are 5 buyers, 5 sellers and 2 competing marketplaces, and let $t_1^b = 4$, $t_2^b = 6$, $t_1^s = 0$, $t_2^s = 2$ and $k_1 = k_2 = 0.5$. Clearly, the traders' evolution of their market selection strategies depends on two factors: the starting point and the fees charged to them. We now choose a starting point (0.8, 0.7, 0.8, 0.7), where traders have higher initial probabilities of choosing marketplace 1. Figure 3.7 then shows the results after we evolve the traders' market selection strategies in the competing marketplaces with different profit fees. The red area is what we call the "lock-in region", which shows when profit fees of marketplace 1 and 2 are within this area, even though marketplace 1 charges a higher profit fee than marketplace 2, traders still converge to marketplace 1. This result is interesting since at this moment, the expensive marketplace can make more profits while still maintaining traders. Note that when the profit fee of marketplace 2 is higher than 60%, marketplace 1 can no longer maintain traders if its profit fee is higher than marketplace 2, i.e. the lock-in region disappears. Furthermore, by running experiments, we also find the existence of the lock-in region when both competing marketplaces charge registration fees, or charge different types of fees. Now we can see that it is possible for traders to converge to the expensive marketplace if currently

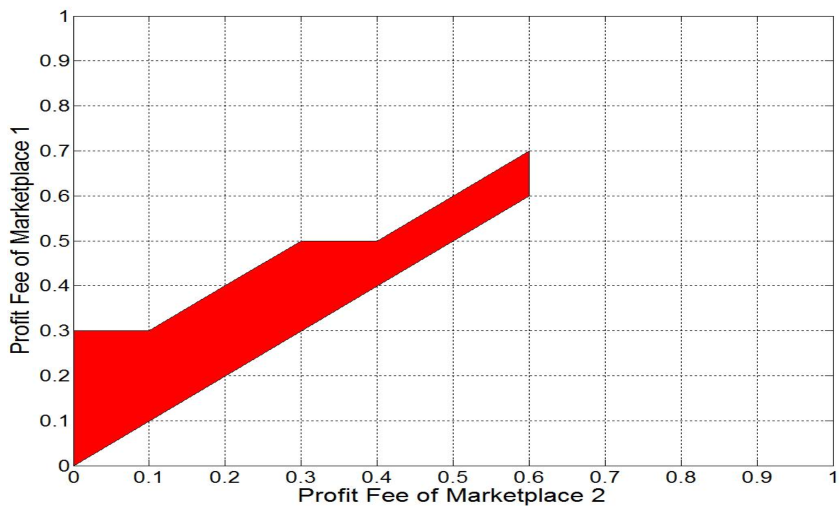

FIGURE 3.7: Lock-in region of marketplace 1 with 5 buyers and 5 sellers.

traders have higher probabilities of choosing this marketplace. This result gives useful insights into a strategy for setting fees in competing marketplaces. Specifically, firstly, a marketplace should lower its fees to attract or maintain traders. After obtaining an advantageous position, the marketplace should then increase its fees higher than its opponents, while still keeping its traders since traders still have higher expected utilities in the expensive marketplace. This so-called bait-and-switch strategy has been adopted by a number of entrants in the CAT competition (Niu et al., 2008a), where initially they charge lower and even no fees to attract traders, and once they have built up a larger market share, they will charge fees to make profits, but still can maintain market share at a good level. While such a strategy is quite intuitive and common in many marketplaces, our analysis provides a more formal justification for it. Furthermore, we can use the strategy as an indication of the level at which the fees should be set.

The effect of the number of traders on the lock-in region:

After obtaining the preliminary conclusion that it is possible for traders to stay in the expensive market, we investigate what factors can affect the size of the lock-in region. In particular, we investigate how the number of traders can affect the size of lock-in region. In the following, we calculate the size of the lock-in region as the sum of the differences of the two marketplaces' discretized profit fees in the lock-in region. For example, the size of the lock-in region in Figure 3.7 is 1.2.[7] From Figure 3.8 we find that, as the number of traders in the competing marketplace environment increases, the size of the lock-in region decreases, which means traders will increasingly select the cheap marketplace. The reason for this is as follows. The traders' choice of marketplaces is determined by their expected utilities, which, in turn, depend on two parts: the gross profit and fees charged to them (see Equation 3.14). From Figure 3.9 we can see that, as the number of traders in the multiple competing marketplaces environment increases, the difference of the traders' gross profits in two marketplaces (i.e. $\Lambda_1^b(\cdot) - \Lambda_2^b(\cdot)$, $\Lambda_1^s(\cdot) - \Lambda_2^s(\cdot)$, see

[7]Since we consider discretized fees, the size of the lock-in region is defined as the sum of differences of the two marketplaces' discretized fees, which is: $(0.3 - 0.0) + (0.3 - 0.1) + (0.4 - 0.2) + (0.5 - 0.3) + (0.5 - 0.4) + (0.6 - 0.5) + (0.7 - 0.6)$.

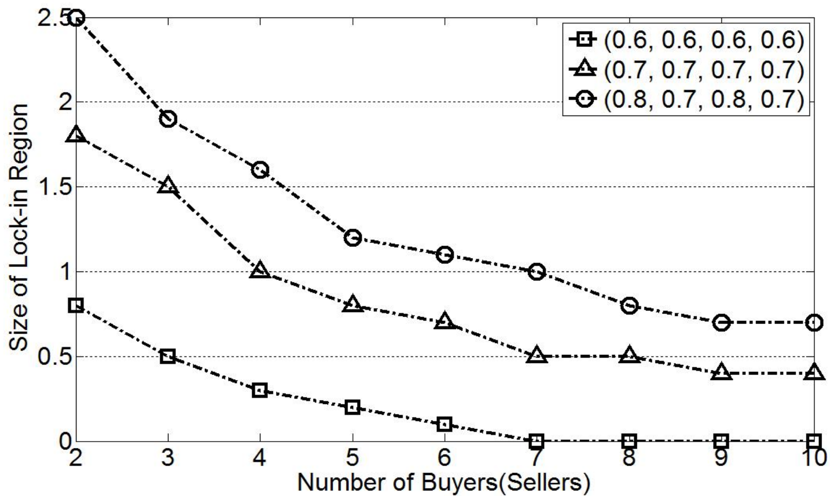

FIGURE 3.8: The size of lock-in region with respect to the number of traders.

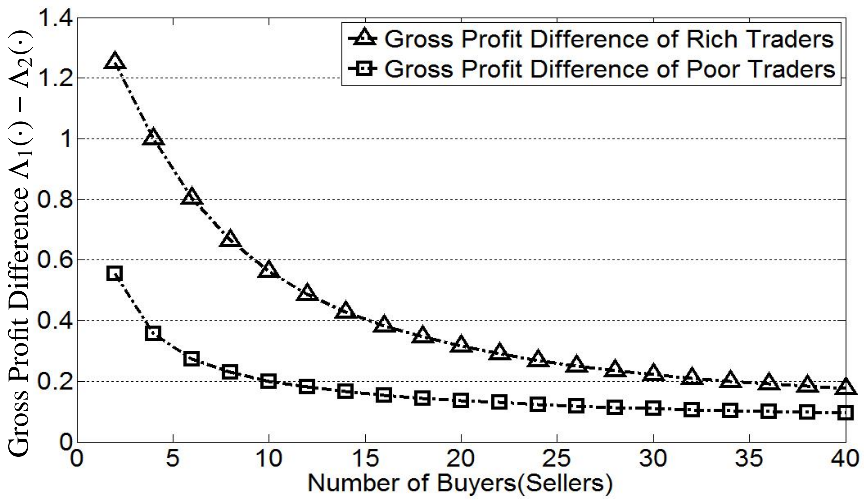

FIGURE 3.9: Gross profit difference in two marketplaces with the mixed strategies of traders (0.8, 0.7, 0.8, 0.7).

Equation 3.13) gradually decreases. This means that the gross profits of traders in two marketplaces gradually become closer to each other. Then the traders' choice of marketplace is mainly determined by the market fees. Thus they will increasingly choose the cheap marketplace. This indicates that, in a multiple competing marketplaces context with a large number of traders, it is difficult for the marketplace to maintain both a high number of traders and high profits.

The effect of the trader types on the lock-in region:

Furthermore, intuitively, competing marketplaces want to attract traders of a rich type since they are more likely to make transactions. Given this, we now analyse what happens when a certain type of trader initially has a bias towards selecting a particular marketplace. First, we consider the rich type's effect on the lock-in region, where we fix the poor traders' probabilities of choosing marketplace 1 to be 0.5, and then change the strategies of the rich traders from 0.55 to 0.95 with step size 0.05. The results are shown in Figure 3.10. From this, we can see

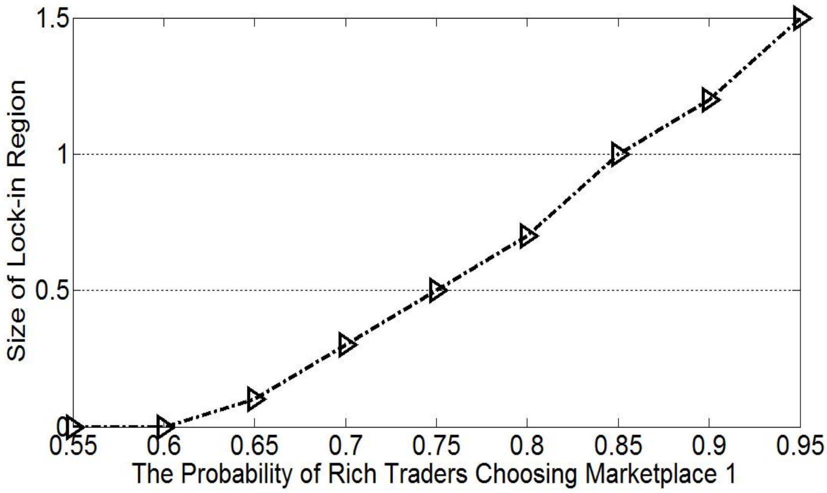

FIGURE 3.10: The size of lock-in region with respect to rich traders' strategy with 5 buyers and 5 sellers.

that, when rich traders have a higher initial probability of selecting marketplace 1, the size of the lock-in region increases. This means that rich traders have a positive effect on the lock-in region. In contrast, if we fix the rich traders' probabilities of choosing marketplace 1 to be 0.5, and then increase the poor traders' probabilities of choosing marketplace 1 starting from 0.05, we find no lock-in regions exist. This is because the surpluses were chosen such that poor traders can make relatively good profits, which means they are not poor enough. Thus we reduce the poor traders' surpluses to enhance their effect on the lock-in region. If we let $t_1^b = 1$, $t_2^b = 8$, $t_1^s = 0$ and $t_2^s = 7$, we get the following result. When the poor traders' probability of choosing market 1 is 0.1, the lock-in region exists and its size is 0.1. When the probability increases to 0.2, the lock-in region disappears. Thus we can see that poor traders have a negative effect on the size of the lock-in region.

The effect of randomisation of market selection strategies on the lock-in region:

So far, we assumed that traders always evolve from their current market selection strategies. Now we analyse how the lock-in region will be affected when some traders are able to explore other marketplaces randomly. We do this because, first of all, traders usually have incomplete information about other traders' market selection decisions. Thus they need to explore and try different marketplaces to obtain more information. For this reason, in the CAT competition, traders have some probability of randomly selecting a marketplace to explore other marketplaces. Secondly, in reality, not all traders are (fully) rational, i.e. they may not always choose the cheapest marketplace. Thus we consider the case where some traders randomly select marketplaces. For example, when the randomisation probability is 10%, then traders will have 90% probability of using their current market selection strategies and 10% probability of selecting each marketplace with equal probability to explore other marketplaces. For this setting, we analyse how the probability of randomisation affects the size of lock-in region. Then the relationship between the randomisation probability and the size of lock-in region is shown in Figure 3.11. We find that when the probability of randomly selecting marketplaces increases,

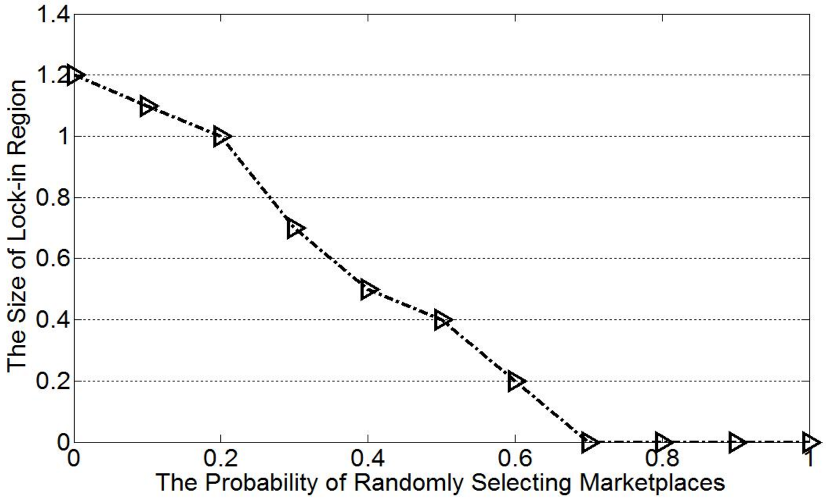

FIGURE 3.11: Relationship between the probability of randomly selecting marketplaces and the size of lock-in region with 5 buyers and 5 sellers.

the size of lock-in region decreases. Furthermore, when the probability of randomly selecting marketplaces is higher than 70%, the lock-in region disappears completely. This means that, as exploration increases, it is more difficult for the competing marketplace to keep traders when charging higher fees even though it initially has a larger market share. Thus in the environment with traders having greater probabilities to explore to search for the cheaper marketplace, the marketplace with a large market share has only a limited advantage.

3.2.3.4 Greater Numbers of Competing Marketplaces

So far we have analysed the market selection strategies of traders in the setting with two competing marketplaces. As stated previously, this analysis is in line with all previous theoretical work which has focused on this canonical case, see Section 2.4.2.1. However, in the real world, it is often the case that more than two marketplaces compete with one another to attract traders and make profits. Now, intuitively, we expect that our results will carry over to this more complex setting. Specifically, when multiple competing marketplaces only charge the same type of fees (i.e. registration or profit fees), we expect that traders will still converge to one marketplace in equilibrium. On the other hand, when multiple competing marketplaces charge different types of fees (i.e. some marketplaces charge profit fees, and others charge registration fees), we believe that traders will either converge to only one marketplace, or only two marketplaces where one charges a profit fee and the other charges a registration fee. To explore these hypotheses, we ran experiments with larger numbers of marketplaces[8]. By so doing, we found that, consistent with the previous analysis, when all marketplaces charge profit fees (or registration fees), traders will converge to one of them in equilibrium[9]. Exactly which one depends on the initial starting

[8]We still assume that there are 5 buyers and 5 sellers, and let $t_1^b = 4$, $t_2^b = 6$, $t_1^s = 0$ and $t_2^s = 2$. For the pricing parameters, we assume that $k_1 = k_2 = 0.5$.

[9]This means that multiple marketplaces cannot co-exist. This conclusion is different from what we observe in practice. We believe that this is because in our model, different marketplaces adopt the same mechanism and have

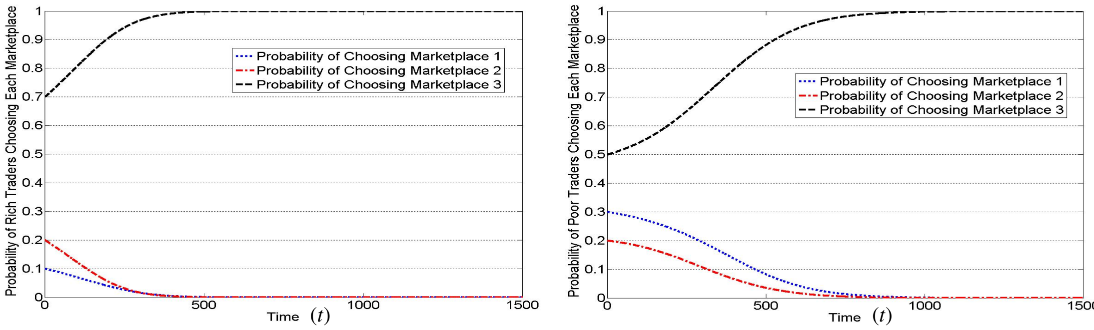

(a) Rich traders.

(b) Poor traders.

FIGURE 3.12: Evolutionary process of traders in the setting with 5 buyers, 5 sellers and 3 competing marketplaces when $q_1 = 10\%$, $q_2 = 20\%$, $q_3 = 30\%$ and $r_1 = r_2 = r_3 = 0$.

point and market fees. Moreover, consistent with the previous analysis, in such experiments, we also found that traders may converge to the expensive marketplace in equilibrium when this marketplace initially has a larger market share. For example, when there are three competing marketplaces where marketplaces 1, 2 and 3 charge 10%, 20% and 30% profit fees respectively, and initially rich traders choose marketplaces 1, 2 and 3 with probabilities 0.1, 0.2 and 0.7 respectively, and poor traders choose marketplaces 1, 2 and 3 with probabilities 0.3, 0.2 and 0.5 respectively, the dynamic changes of traders' probabilities of choosing marketplaces are shown in Figure 3.12. From the evolutionary results, we can see that eventually traders choose one marketplace with 100% probability, i.e. converge to one marketplace. Specifically, in this case, traders converge to marketplace 3 which is the most expensive since initially this marketplace has a larger market share. Furthermore, when some of them charge registration fees and others charge profit fees, we found that traders either converge to one marketplace in equilibrium or the rich traders converge to the marketplace which charges a registration fee and the poor traders converge to the marketplace which charges a profit fee (multiple competing marketplaces which charge the same type of fees do not co-exist in equilibrium and only one of them can survive). For example, when there are three competing marketplaces where marketplaces 1 and 2 charge 40% and 50% profit fees respectively and marketplace 3 charges a 0.8 registration fee, and initially rich traders choose marketplaces 1, 2 and 3 with probabilities 0.15, 0.2 and 0.65 respectively, and poor traders choose marketplaces 1, 2 and 3 with probabilities 0.35, 0.25 and 0.4 respectively, the dynamic changes of traders' probabilities of choosing marketplaces are shown in Figure 3.13. From this figure, we can see that eventually rich traders converge to marketplace 3 charging a registration fee, and poor traders converge to marketplace 1 charging a profit fee. We also can see that marketplaces 1 and 2 charging the same types of fees cannot co-exist, and only marketplace 1 survives.

identical goods, and thus cannot provide enough diversity for traders to select different marketplaces.

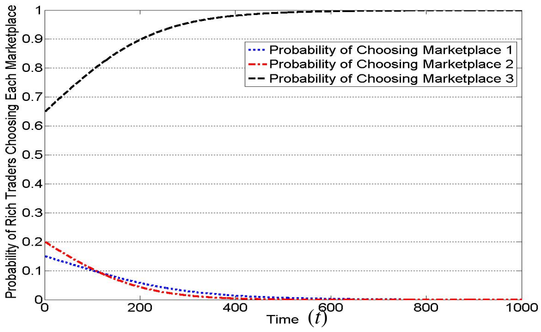

(a) Rich traders.

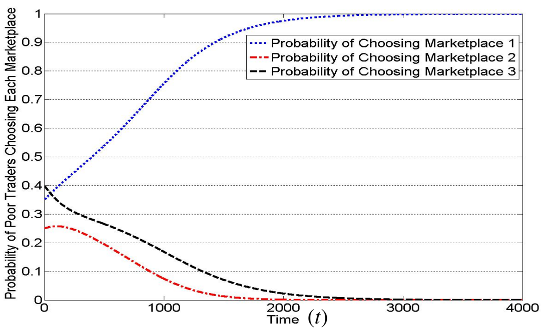

(b) Poor traders.

FIGURE 3.13: Evolutionary process of traders in the setting with 5 buyers, 5 sellers and 3 competing marketplaces when $q_1 = 45\%$, $q_2 = 50\%$, $q_3 = 0$, $r_1 = r_2 = 0$ and $r_3 = 0.8$.

3.3 Equilibrium Analysis of Charging Strategies

In the previous subsection, we analysed the traders' equilibrium strategies of market selection for a given fee system. Now, given the insights from this analysis, we analyse how marketplaces should set fees to make profits in equilibrium. In the following, we analyse this problem through two different approaches. In the first, we investigate the equilibrium charging strategies from a static analysis. This approach is based on the assumption that marketplaces set their fees once at the beginning and so the charging strategies are not affected by the changes in the traders' market selection strategies. In the second approach, we address this limitation by modelling the game as a two-stage game where the strategies of the traders and the marketplaces are affected by each other.

3.3.1 Static Analysis

In this section, we derive equilibrium charging strategies for the marketplaces through a static analysis, where we assume that the marketplaces' charging strategies are not affected by dynamic changes in the traders' market selection strategies (although they are affected by the traders' equilibrium market selection strategies, which is described in the following). Specifically, in this analysis, we first derive the marketplaces' expected utilities corresponding to each possible fee system, which are dependent on the traders' equilibrium market selection strategies in this fee system and the probability of traders converging to that equilibrium. Specifically, we use EGT to analyse how traders evolve their market selection strategies in the given fee system, and approximate the probability of traders converging to a specific equilibrium by approximating the size of basin of attraction of that equilibrium under the assumption that each starting point is equally likely selected by traders. After calculating the marketplaces' expected utilities for each possible fee system, we obtain the payoff table[10]. Finally, we analyse the equilibrium

[10]It is a matrix, where the first column(row) represents the corresponding marketplace's fee, and the cell represents the marketplaces' utilities, which correspond to marketplaces' fees.

charging strategies according to the payoff table. We can see that in this analysis, how the marketplaces' expected utilities are calculated is a key issue. Therefore, in the following, we first derive equations to calculate these expected utilities. Based on this, we will analyse how marketplaces set fees in equilibrium in two different cases where marketplaces charge the same type of fees and different types of fees.

3.3.1.1 Expected Utilities of Marketplaces

In this section, we describe how to calculate the marketplaces' expected utilities for a given market fee system $\bar{P}$. Intuitively, we can see that the marketplaces' expected utilities not only depend on the current fee system $\bar{P}$, but also on the traders' market selection strategy profiles $\bar{\omega}^b(\bar{P}) = \langle \omega^b(t_1^b, m, \bar{P}), \omega^b(t_2^b, m, \bar{P}) \rangle$ and $\bar{\omega}^s(\bar{P}) = \langle \omega^s(t_1^s, m, \bar{P}), \omega^s(t_2^s, m, \bar{P}) \rangle$, which is conditional on the fee system. In order to derive marketplace m's expected utility, we first calculate the probability that there are exactly τ_1^b poor buyers and τ_2^b rich buyers choosing m:

$$\varrho_m^b(\tau_1^b, \tau_2^b) = \binom{B}{\tau_1^b, \tau_2^b, B - \tau_1^b - \tau_2^b} * \left(\frac{\omega^b(t_1^b, m, \bar{P})}{2}\right)^{\tau_1^b} * \left(\frac{\omega^b(t_2^b, m, \bar{P})}{2}\right)^{\tau_2^b} * \left(1 - \frac{\omega^b(t_1^b, m, \bar{P})}{2} - \frac{\omega^b(t_2^b, m, \bar{P})}{2}\right)^{B-\tau_1^b-\tau_2^b} \quad (3.26)$$

where $\binom{B}{\tau_1^b, \tau_2^b, B-\tau_1^b-\tau_2^b}$ is the multinomial coefficient and $\frac{\omega^b(t_1^b, m, \bar{P})}{2}$ is the probability that a buyer is poor and chooses marketplace m. Similarly, we get the probability that there are exactly τ_1^s rich sellers and τ_2^s poor sellers in m:

$$\varrho_m^s(\tau_1^s, \tau_2^s) = \binom{S}{\tau_1^s, \tau_2^s, S - \tau_1^s - \tau_2^s} * \left(\frac{\omega^s(t_1^s, m, \bar{P})}{2}\right)^{\tau_1^s} * \left(\frac{\omega^s(t_2^s, m, \bar{P})}{2}\right)^{\tau_2^s} * \left(1 - \frac{\omega^s(t_1^s, m, \bar{P})}{2} - \frac{\omega^s(t_2^s, m, \bar{P})}{2}\right)^{S-\tau_1^s-\tau_2^s} \quad (3.27)$$

Furthermore, marketplace m's expected utility in the fee system $\bar{P}$ given its pricing parameter k_m when there are exactly τ_1^b poor buyers, τ_2^b rich buyers, τ_1^s rich sellers and τ_2^s poor sellers in this marketplace is calculated by:

$$\tilde{U}_m(\bar{P}, k_m, \tau_1^b, \tau_2^b, \tau_1^s, \tau_2^s) = (\tau_1^b + \tau_2^b + \tau_1^s + \tau_2^s) * r_m + (\Lambda^b + \Lambda^s) * q_m \quad (3.28)$$

where Λ^b, Λ^s are the buyers' and the sellers' share of the trading surplus respectively when τ_1^b poor buyers, τ_2^b rich buyers, τ_1^s rich sellers and τ_2^s poor sellers are matched according to the equilibrium matching policy[11]. At this moment, we can get the marketplace's expected utility given the fee system and the traders' market selection strategy profiles:

$$\tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P})) = \sum_{\tau_1^b=0}^{B} \sum_{\tau_2^b=0}^{B-\tau_1^b} \sum_{\tau_1^s=0}^{S} \sum_{\tau_2^s=0}^{S-\tau_1^s} *\varrho_m^b(\tau_1^b, \tau_2^b) * \varrho_m^s(\tau_1^s, \tau_2^s) * \tilde{U}_m(\bar{P}, k_m, \tau_1^b, \tau_2^b, \tau_1^s, \tau_2^s) \quad (3.29)$$

Now given the fee system $\bar{P}$, we need to calculate the marketplace's expected utility at the point where all traders use equilibrium market selection strategies, which are conditional on

[11]This can be easily calculated. For example, when there are 2 rich buyers, 3 poor buyers, 3 rich sellers and 2 poor buyers in marketplace m, $\Lambda^b = (max(t_2^b - t_1^s, 0) * 2 + max(t_1^b - t_1^s, 0) + max(t_1^b - t_2^s, 0) * 2) * k_m$ and $\Lambda^s = (max(t_2^b - t_1^s, 0) * 2 + max(t_1^b - t_1^s, 0) + max(t_1^b - t_2^s, 0) * 2) * (1 - k_m)$.

this fee system. As we discussed previously, there can exist multiple BNEs. In such cases, the marketplace's expected utility depends on which BNE strategy the traders will choose and the probability of choosing this BNE strategy. Given a fee system, we have used EGT to analyse how traders choose BNE strategies in Section 3.2.3.2. Similarly, here, we use EGT to find which BNE strategies traders will choose and with what probability. Recall that in EGT, we use the replicator dynamics to show the trajectories and how they converge to an equilibrium. The size of the basin, where all trajectories converge to a particular equilibrium, can be used to indicate the probability of traders converging to that equilibrium (see Section 3.2.3.1). However, there are infinitely many possible starting points, which means that we have to approximate the size of basins. In this work, we do this by discretizing the starting points. Specifically, we calculate the size of basin of attraction by discretizing the mixed strategy of each type from 0.01 to 0.99 with step size 0.049, which gives $21^4 = 194481$ different starting points[12]. Note that if we use even more points, we can estimate the probability of traders' convergence to each equilibrium more accurately.

Now that we know, given a fee system $\bar{P}$, what BNE strategies traders will choose and with what probabilities, we are able to calculate the expected utility for a marketplace. Specifically, given that there are X possible BNEs, we use $\langle x_1, x_2, ..., x_X\rangle$ to represent the probabilities of traders converging to these BNEs. Then marketplace m's expected utility in the fee system $\bar{P}$ is:

$$\tilde{U}_m(\bar{P}) = \sum_{z=1}^{X} x_z * \tilde{U}_m(\bar{P}, k_m, \langle\omega^{zb}(t_1^b, m, \bar{P}), \omega^{zb}(t_2^b, m, \bar{P})\rangle, \langle\omega^{zs}(t_1^s, m, \bar{P}), \omega^{zs}(t_2^s, m, \bar{P})\rangle) \quad (3.30)$$

where $\omega^{zb}(t_1^b, m, \bar{P})$, $\omega^{zb}(t_2^b, m, \bar{P})$, $\omega^{zs}(t_1^s, m, \bar{P})$ and $\omega^{zs}(t_2^s, m, \bar{P})$ denote the z-th BNE market selection strategies.

Now we have derived equations to calculate the marketplace's expected utility given the fee system $\bar{P}$. Based on this, we can analyse the equilibrium charging strategy for marketplaces. As we know, the range of possible fees is continuous, which results in infinitely many possible fee systems. It is too complicated to analyse the game with an infinite strategy space. However, in Wellman (2006), researchers claim that for this kind of game, it is useful to approximate the game by restricting the strategy space, and results from the restricted strategy space still provide insights into the original game. Similarly, in this work, in order to obtain tractable results, we also restrict the fee space by discretizing these fees. Then we calculate the marketplaces' expected utilities corresponding to these fees, and generate the payoff table for marketplaces, by which we can analyse the equilibrium fee system.

3.3.1.2 Experiment Results

After deriving equations to calculate the marketplaces' expected utilities and describing the analysis process, we now analyse the equilibrium charging strategies for marketplaces. We first

[12]The discretization is not from 0 to 1 since when the mixed strategy is 0 or 1, it constitutes an equilibrium of replicator dynamics, but may be not a Nash equilibrium.

consider the case that only profit fees can be charged to traders. We still assume that there are 5 buyers, 5 sellers and 2 competing marketplaces, and let $t_1^b = 4$, $t_2^b = 6$, $t_1^s = 0$ and $t_2^s = 2$. For the pricing parameters, we assume that $k_1 = k_2 = 0.5$.[13] Furthermore, we discretize profit fees from 0 to 1 with step size 0.1. Therefore, each marketplace can choose from 11 different profit fees. For two competing marketplaces, there are $11^2 = 121$ different fee systems. For each of these combinations, we use EGT to obtain the basin of attraction to each BNE of the market selection strategies. Then by approximating the size of each basin, we get the probability of traders choosing each BNE, which is shown in Figure 3.14. Then using Equation 3.30, we calculate the marketplaces' expected profits. The results are shown in Table 3.1. From this table, by using Gambit (`http://gambit.sourceforge.net`), we find that both marketplaces charging 30% profit fee constitutes a unique pure Nash equilibrium (NEQ) fee system. In this equilibrium, both competing marketplaces charge non-zero profit fees and therefore make positive profits. Furthermore, we also analyse the case that both competing marketplaces only charge registration fees, and find that both marketplaces charging a 0.1 registration fee constitutes a unique NEQ.

Now we will consider the case where different competing marketplaces charge different types of fees: marketplace 1 charges only profit fees, and marketplace 2 charge only registration fees. We discretize registration and profit fees from 0 to 1 with step size 0.1. Then there are again 121 different fee systems. By exploring the traders' market selection strategies under all possible fee systems, we obtain the probabilities of traders converging to each BNE, which are shown in Figure 3.15. Note that, as we discussed in Section 3.2.3.2, when different types of fees are charged, traders may converge to different marketplaces in the equilibrium. Figure 3.15(c) shows in which fee systems, traders may converge to different marketplaces. After estimating the probabilities of traders' convergence to each BNE, we then calculate the marketplaces' expected utilities using Equation 3.30. The marketplaces' expected utilities are shown in Table 3.2, from which we can see that in this case, marketplace 1 charging 30% profit fee, and marketplace 2 charging 0.5 registration fee constitutes the unique NEQ fee system.

[13] In this situation, rich(poor) buyer and rich(poor) seller have the same behaviour.

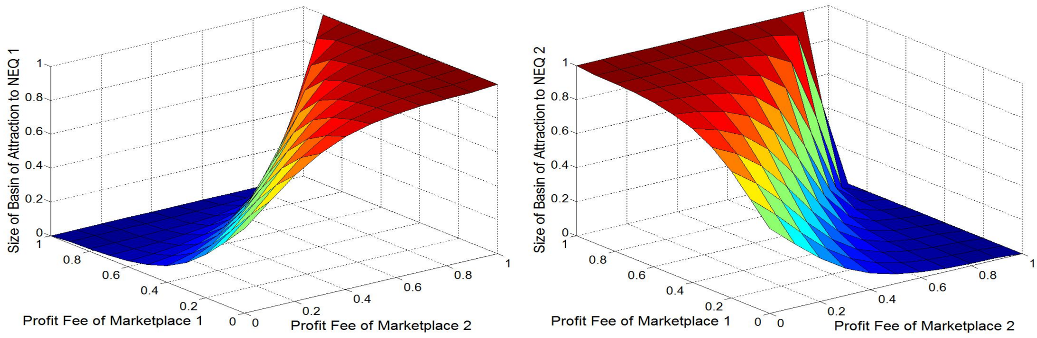

(a) Size of basin of attraction to BNE 1 (marketplace 1). (b) Size of basin of attraction to BNE 2 (marketplace 2).

FIGURE 3.14: Sizes of basins of attraction when both competing marketplaces only charge profit fees.

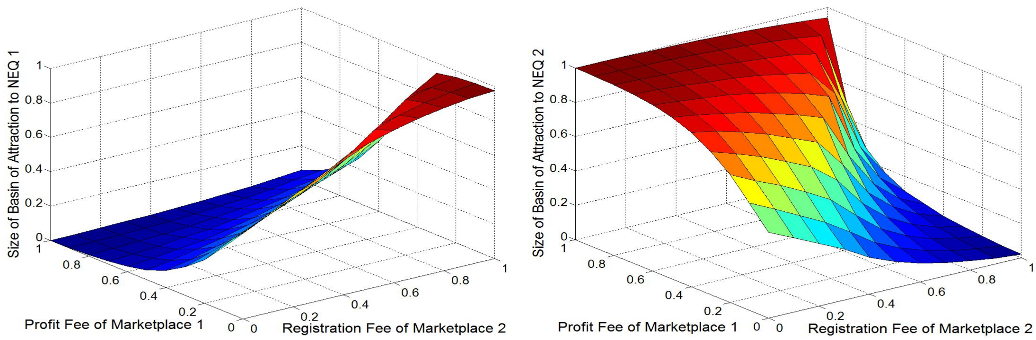

(a) Size of basin of attraction to BNE 1 (marketplace 1). (b) Size of basin of attraction to BNE 2 (marketplace 2).

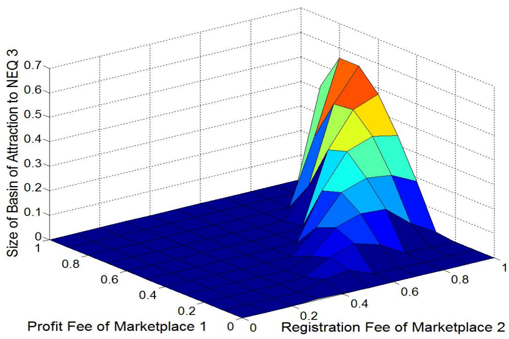

(c) Size of basin of attraction to BNE 3.

FIGURE 3.15: Sizes of basins of attraction when competing marketplaces charge different types of fees.

	0.0	**0.1**	**0.2**	***0.3***	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
0.0	0.00,0.00	0.00,0.75	0.00,1.02	0.00,0.97	0.00,0.78	0.00,0.58	0.00,0.42	0.00,0.29	0.00,0.22	0.00,0.20	0.00,0.00
0.1	0.75,0.00	1.00,1.00	1.28,1.44	1.54,1.39	1.73,1.10	1.85,0.77	1.91,0.52	1.95,0.34	1.97,0.24	1.98,0.20	2.00,0.00
0.2	1.02,0.00	1.44,1.28	2.00,2.00	2.63,2.05	3.18,1.64	3.56,1.11	3.77,0.70	3.88,0.43	3.93,0.27	3.95,0.22	4.00,0.00
0.3	0.97,0.00	1.39,1.54	2.05,2.63	**3.00,3.00**	4.07,2.58	4.97,1.72	5.50,1.01	5.76,0.56	5.88,0.32	5.93,0.22	6.00,0.00
0.4	0.78,0.00	1.10,1.73	1.64,3.18	2.58,4.07	4.00,4.00	5.66,2.93	6.90,1.64	7.54,0.81	7.81,0.38	7.90,0.22	8.00,0.00
0.5	0.58,0.00	0.77,1.85	1.11,3.56	1.72,4.97	2.93,5.66	5.00,5.00	7.45,3.06	9.03,1.36	9.65,0.56	9.86,0.25	10.00,0.00
0.6	0.42,0.00	0.52,1.91	0.70,3.77	1.01,5.50	1.64,6.90	3.06,7.45	6.00,6.00	9.54,2.87	11.30,0.93	11.80,0.31	12.00,0.00
0.7	0.29,0.00	0.34,1.95	0.43,3.88	0.56,5.76	0.81,7.54	1.36,9.03	2.87,9.54	7.00,7.00	12.08,2.19	13.66,0.43	14.00,0.00
0.8	0.22,0.00	0.24,1.97	0.27,3.93	0.32,5.88	0.38,7.81	0.56,9.65	0.93,11.30	2.19,12.08	8.00,8.00	15.07,1.04	16.00,0.00
0.9	0.20,0.00	0.20,1.98	0.22,3.95	0.22,5.93	0.22,7.90	0.25,9.86	0.31,11.80	0.43,13.66	1.04,15.07	9.00,9.00	18.00,0.00
1.0	0.00,0.00	0.00,2.00	0.00,4.00	0.00,6.00	0.00,8.00	0.00,10.0	0.00,12.00	0.00,14.00	0.00,16.00	0.00,18.00	10.00,10.00

Table 3.1: Expected utilities of marketplace 1 and marketplace 2. The first column is the profit fee of marketplace 1 and the first row is the profit fee of marketplace 2. The first element in each cell is marketplace 1's expected utility, and the second is marketplace 2's expected utility. Bold italic fees constitute a NEQ fee system.

	0.0	**0.1**	**0.2**	**0.3**	**0.4**	***0.5***	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
0.0	0.00,0.00	0.00,0.43	0.00,0.73	0.00,0.90	0.00,0.86	0.00,0.77	0.00,0.67	0.00,0.54	0.00,0.43	0.00,0.31	0.00,0.23
0.1	0.75,0.00	0.89,0.56	1.03,0.97	1.18,1.22	1.32,1.32	1.49,1.20	1.69,0.92	1.79,0.73	1.86,0.54	1.91,0.39	1.95,0.27
0.2	1.02,0.00	1.31,0.67	1.55,1.22	1.84,1.62	2.15,1.85	2.48,1.81	2.83,1.59	3.22,1.22	3.62,0.75	3.76,0.54	3.86,0.36
0.3	1.00,0.00	1.29,0.79	1.59,1.47	1.97,2.01	2.37,2.42	**2.88,2.55**	3.44,2.33	4.00,1.99	4.49,1.64	5.20,1.01	5.67,0.53
0.4	0.80,0.00	1.03,0.87	1.34,1.66	1.72,2.36	2.19,2.90	2.66,3.34	3.35,3.30	4.17,2.89	4.91,2.48	5.62,2.03	6.32,1.53
0.5	0.57,0.00	0.79,0.92	1.07,1.79	1.36,2.59	1.77,3.29	2.20,3.90	2.72,4.37	3.71,3.99	4.62,3.49	5.56,2.99	6.38,2.56
0.6	0.38,0.00	0.54,0.96	0.76,1.87	0.92,2.77	1.28,3.57	1.66,4.31	2.12,4.94	2.64,5.44	3.92,4.64	4.99,4.03	5.97,3.54
0.7	0.22,0.00	0.32,0.98	0.48,1.93	0.70,2.85	0.85,3.76	1.15,4.59	1.53,5.35	1.93,6.03	2.73,6.14	4.09,5.18	5.19,4.55
0.8	0.11,0.00	0.18,0.99	0.29,1.96	0.37,2.93	0.54,3.86	0.77,4.76	1.01,5.62	1.26,6.45	1.71,7.14	2.86,6.66	4.35,5.54
0.9	0.04,0.00	0.09,1.00	0.16,1.98	0.20,2.97	0.29,3.94	0.45,4.88	0.61,5.80	0.81,6.69	1.10,7.51	1.48,8.26	3.24,6.92
1.0	0.00,0.00	0.00,1.00	0.04,2.00	0.10,2.99	0.14,3.97	0.22,4.95	0.36,5.89	0.50,6.83	0.68,7.73	0.90,8.60	1.18,9.41

Table 3.2: Expected utilities of marketplace 1 and marketplace 2. The first column is the profit fee of marketplace 1 and the first row is the registration fee of marketplace 2. The first element in each cell is marketplace 1's expected utility, and the second is marketplace 2's expected utility. Bold italic fees constitute a NEQ fee system.

3.3.2 Co-Evolutionary Analysis

In the above, we have investigated the marketplaces' equilibrium charging strategies through a static analysis, which is restricted to the assumption that their charging strategies do not change in response to the traders' market selection strategies. In this section, we address this limitation by modelling the game as a two-stage game where, in the first stage, marketplaces publish their fee structures according to their charging strategies and then, in the second stage, traders select marketplaces according to their market selection strategies, which are conditional on the fee system from the first stage. Given this complicated setting of a two-stage game with incomplete information about traders' types, it is difficult to use traditional game-theoretic methods to analyse equilibrium charging strategies. Intuitively, we can see that the traders' market selection strategies and the marketplaces' charging strategies will affect each other. Hence, we use a co-evolutionary approach to analyse this problem. This approach can capture the dynamic process of how marketplaces evolve their charging strategies to converge to equilibrium while taking into account the dynamic changes of the traders' market selection strategies. In the following, before we perform the co-evolutionary analysis, we first describe the co-evolutionary process in more detail.

3.3.2.1 The Co-Evolutionary Process

In the co-evolutionary process, both the competing marketplaces and the traders dynamically learn to adapt their strategies to maximise their own expected utilities. This learning process is repeated until both traders and marketplaces do not change their strategies. At this moment, an equilibrium is reached. In each learning round (i.e. a co-evolutionary step), traders update their expected utilities before they evolve their market selection strategies. Now in order to calculate the expected utilities, they require information about the local type distribution (i.e. the type distribution of traders in a specific marketplace, see Section 3.2.1) of other traders. While in Section 3.2.1 the local type distribution was calculated for a given fee system, the strategies of the marketplaces are mixed (i.e. each fee system is selected by marketplaces with a certain probability) and therefore in this case we calculate the local type distribution taking into account the mixed marketplace strategies. Specifically, these local type distributions depend on the traders' market selection strategies (which are conditional on each fee system) and the marketplaces' charging strategies (which determine the probability of each possible fee system being selected). This is important since it creates a link between the strategy composition of the marketplaces and its effect on the expected utility of the traders, enabling co-evolution to occur[14].

Now we describe the co-evolutionary process in detail, which is depicted in Figure 3.16. First, we initialise the marketplaces' charging strategies and the traders' market selection strategies. Then we calculate the initial local type distributions of buyers and sellers (see Equation 3.32 in

[14] An alternative approach is to keep the local distribution conditional on the fees, but then the population dynamics of the charging strategies will have no effect on the traders' expected utilities.

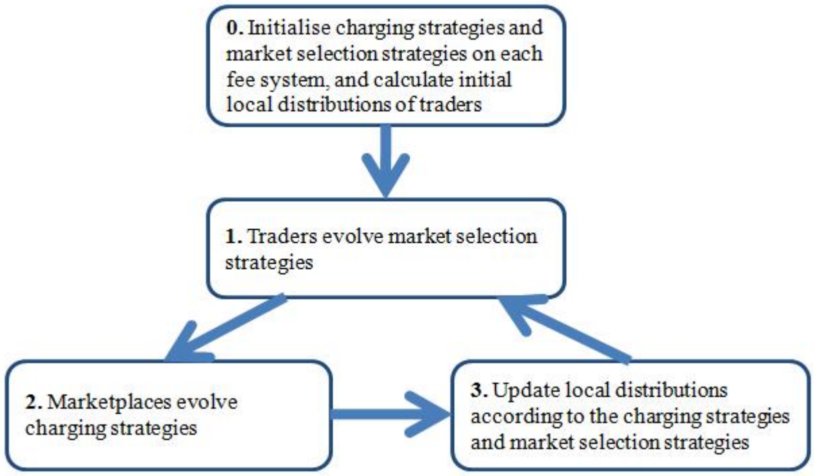

FIGURE 3.16: The co-evolutionary process.

the following subsection). From the initial local distributions, traders calculate their expected utilities, and then evolve their market selection strategies. After traders evolve their market selection strategies, the marketplaces calculate their expected utilities (which depend on the traders' mixed market selection strategies and the marketplaces' charging strategies) and then evolve their charging strategies. After the marketplaces evolve their charging strategies, we update the local distributions of the traders, and then enter the next co-evolutionary step. This co-evolutionary process continues until all dynamic changes of traders' market selection strategies and marketplaces' charging strategies become zero. At this point, an equilibrium is reached.

In the following, before giving the experimental analysis in detail, we first need to derive equations to calculate the expected utilities of traders and marketplaces, and give the replicator dynamics equations.

3.3.2.2 Expected Utilities of Traders and Marketplaces

As discussed above, in the co-evolutionary process, a trader's expected utility depends on the local type distributions of the other traders. In Subsection 3.2.1, we calculated the local type distributions for a given fee system. However, since the charging strategies of the marketplaces are mixed, here we consider the traders' local type distributions under all allowable fee systems. These new local type distributions are derived in the following way. First, the probability that the type of a buyer is less than θ^b in marketplace m is:

$$H_m^b(\theta^b) = \sum_{\bar{P} \in \mathcal{P}^M} \mu(\bar{P}) * H_m^b(\theta^b|\bar{P}) = \sum_{\bar{P} \in \mathcal{P}^M} \mu(\bar{P}) * \int_{\underline{l}}^{\theta^b} f^b(x) * \omega^b(x, m, \bar{P}) dx \tag{3.31}$$

where $\mu(\bar{P}) = \prod_{m \in \mathcal{M}} \mu_m(p_m)$ is the probability of the fee system $\bar{P} = \langle p_1, ..., p_M \rangle$ appearing, and $H_m^b(\theta^b|\bar{P})$ is the probability that the type of a buyer is less than θ^b in marketplace m for a given fee system $\bar{P}$ (see Equation 3.4). Then, by normalising the above equation, we obtain a proper local type distribution of the buyers in marketplace m:

$$G_m^b(\theta^b) = \frac{H_m^b(\theta^b)}{H_m^b(\bar{l})} \tag{3.32}$$

The local probability density function of buyer types is:

$$g_m^b(\theta^b) = \frac{\sum_{\bar{P}\in\mathcal{P}^M} \mu(\bar{P}) * f^b(\theta^b) * \omega^b(\theta^b, m, \bar{P})}{H_m^b(\bar{l})} \tag{3.33}$$

The equations for sellers can be calculated in the same way. All other equations to calculate the traders' expected utilities are the same as before except that in these equations we need to replace $H_m^b(\theta^b|\bar{P})$, $G_m^b(\theta^b|\bar{P})$ and $g_m^b(\theta^b|\bar{P})$ (which are conditional on a specific fee system $\bar{P}$) by $H_m^b(\theta^b)$, $G_m^b(\theta^b)$ and $g_m^b(\theta^b)$ (which are under all possible fee systems) respectively.

In addition to the traders' expected utilities, in this two-stage game, we also need to calculate the expected utility of each marketplace. In Section 3.3.1.1, we have derived the marketplace's expected utility given a particular fee system $\bar{P}$ and the traders' market selection strategy profiles: $\bar{\omega}^b(\bar{P})$ and $\bar{\omega}^s(\bar{P})$. This calculation assumes that each marketplace adopts a pure charging strategy, i.e. $\mu_m(p_m) = 1$ $(m = 1, ..., M)$. Here, we need to calculate the expected utility of the marketplace adopting a mixed charging strategy. Intuitively, a marketplace's expected utility not only depends on its own charging strategy, but also on the charging strategies of the other marketplaces and the traders' mixed market selection strategies (which are conditional on the fee systems announced in the first stage). In the following, we calculate the expected utility of marketplace m given a charging strategy profile $\bar{\mu}$ and the market selection strategy profiles of traders $\bar{\omega}^b(\cdot)$ and $\bar{\omega}^s(\cdot)$. In the first step, we use Equation 3.29 to calculate $\tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}))$, which is marketplace m's expected utility given the fee system $\bar{P}$, pricing parameter k_m, the buyers' strategy profile $\bar{\omega}^b(\bar{P}) = \langle \omega^b(t_1^b, m, \bar{P}), \omega^b(t_2^b, m, \bar{P}) \rangle$ and the sellers' strategy profile $\bar{\omega}^s(\bar{P}) = \langle \omega^s(t_1^s, m, \bar{P}), \omega^s(t_2^s, m, \bar{P}) \rangle$. Then marketplace m's expected utility with a (mixed) charging strategy profile $\bar{\mu}$ is:

$$\tilde{U}_m(\bar{\mu}) = \sum_{\bar{P}\in\mathcal{P}^M} \mu(\bar{P}) * \tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P})) \tag{3.34}$$

where $\mu(\bar{P}) = \prod_{m\in\mathcal{M}} \mu_m(p_m)$.

3.3.2.3 Replicator Dynamics

We now describe the replicator dynamics equations for traders and marketplaces respectively for the two-stage game. In addition to adding the replicator dynamics equations for the marketplaces, the two-stage game also requires a considerable increase in the number of equations for the traders. This is because, while in Section 3.2.3, the equations were for a given fee system, these are now conditional on the fee system. For *each possible fee system*, a different population evolves the strategy of each trader type (there are 4 populations: rich buyer, poor buyer, rich seller and poor seller). Specifically, when the fee system is $\bar{P}$, the replicator dynamics equations for each population of traders are given by Equations 3.22, 3.23, 3.24 and 3.25. We can see that, for each fee system, there are $4 \times M$ replicator dynamics equations. Since there are $|\mathcal{P}^M|$ different fee systems, in total, there are $|\mathcal{P}^M| \times 4 \times M$ replicator equations for traders.

Now we describe replicator dynamics equations for the marketplaces. Since there are $|\mathcal{P}|$ allowable fee structures, for each marketplace, there are $|\mathcal{P}|$ replicator dynamics equations for its charging strategy. In total there are $M \times |\mathcal{P}|$ replicator dynamics equations for marketplaces[15]. Specifically, marketplace m's replicator dynamics equation for fee structure p'_m is as follows:

$$\dot{\mu}_m(p'_m) = \frac{d\mu_m(p'_m)}{dt} = \left(\tilde{U}_m(p'_m, \mu_{-m}(\cdot)) - \tilde{U}_m(\bar{\mu})\right) * \mu_m(p'_m) \tag{3.35}$$

where $\dot{\mu}_m(p'_m)$ describes how marketplace m changes its probability of choosing fee structure p'_m, $\tilde{U}_m(\bar{\mu})$ is m's overall expected utility as derived in Section 3.3.2.2, and $\tilde{U}_m(p'_m, \mu_{-m}(\cdot))$ is m's expected utility of choosing fee structure p'_m given the other marketplaces' charging strategy profile $\mu_{-m}(\cdot)$:

$$\tilde{U}_m(p'_m, \mu_{-m}(\cdot)) = \sum_{\bar{P} \in \mathcal{P}^M:\, p_m = p'_m} \prod_{l \in \mathcal{M} \setminus \{m\}} \mu_l(p_l) * \tilde{U}_m(\bar{P}, k_m, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}))) \tag{3.36}$$

3.3.2.4 Experimental Results

After describing the co-evolutionary process and the replicator dynamics equations, we are ready to analyse how marketplaces evolve their charging strategies over time. Specifically, in the following analysis, we still discretize the profit and registration fees from 0 to 1 with step size 0.1. We also assume that there are 2 competing marketplaces, 5 buyers and 5 sellers, and let $t_1^b = 4$, $t_2^b = 6$, $t_1^s = 0$, $t_2^s = 2$. Furthermore, for the pricing parameters, we assume that $k_1 = k_2 = 0.5$.

Two identical marketplaces initially having the same charging strategy:

First, we consider that both marketplaces only charge profit fees. Then there are 11 possible fee structures[16] for each marketplace, which implies that there are 968 replicator dynamics equations for the traders and 22 replicator dynamics equations for the two competing marketplaces. We assume that initially both marketplaces are identical. That is, they have the same probabilities of choosing each fee structure, and for each fee system, the initial probability of traders choosing marketplace 1(2) is the ratio of profit fee of marketplace 2(1) to the sum of profit fees of both competing marketplaces (this means that traders have higher initial probabilities of choosing the cheaper marketplace). Then the initial probability of traders choosing each marketplace under all possible fee systems is equal, i.e. 0.5. From this setting, we evolve the charging

[15] Here we consider that different marketplaces are from different populations since different marketplaces may adopt different types of fees, or even though they use the same type of fees, different marketplace may set their initial charging strategies differently.

[16] They are (0,0), (0,0.1), (0,0.2), (0,0.3), (0,0.4), (0,0.5), (0,0.6), (0,0.7), (0,0.8), (0,0.9), (0,1.0).

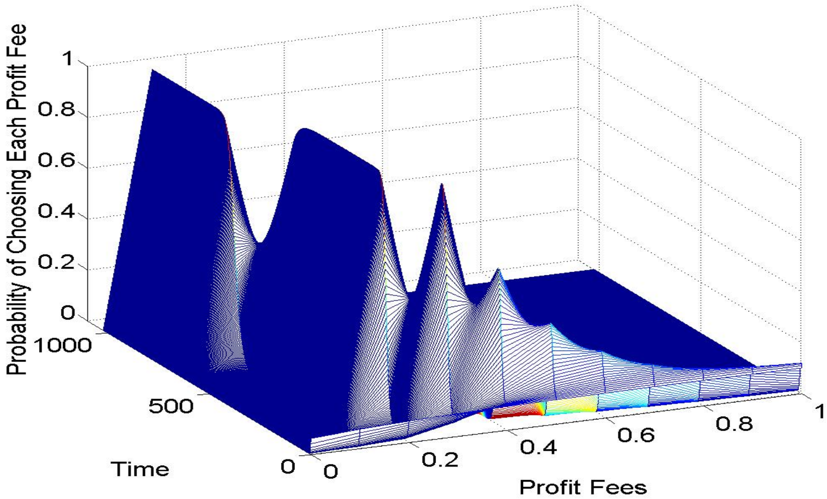

FIGURE 3.17: Evolutionary process of charging strategies of marketplace 1 and 2 when they have identical initial charging strategies.

strategies of marketplaces and the market selection strategies of the traders. The evolutionary process of charging strategies is shown in Figure 3.17 where the x-axis is the possible profit fees the marketplace can charge, the y-axis is the evolutionary time, and the z-axis is the probability of the marketplace choosing each profit fee during the evolutionary process. Note that, in this case, the evolutionary processes of two initially identical marketplaces are still identical. From the figure, we can see that during the evolutionary process, both marketplaces gradually set low fees with higher probability, and in equilibrium, both marketplaces set a 10% profit fee with 100% probability. This is because two identical marketplaces have to undercut each other by decreasing fees to attract traders. Eventually, they converge to a pure strategy where both marketplaces charge a 10% profit fee, which is the minimum allowed profit fee which can guarantee positive profit for the marketplaces[17]. In addition, for the traders' evolutionary process, we look at the traders' probability of choosing marketplace 1 considering all possible fee systems. In this case, the probability of traders choosing each marketplace is unchanged, which is 0.5. This shows that in a highly competitive environment, competing marketplaces have to charge the lowest fees, and even no fees, in order to keep traders. This is different from the static analysis in Section 3.3.1, where in equilibrium both competing marketplaces charge a 30% profit fee. The reason is that in the co-evolutionary analysis, competing marketplaces respond to the dynamic changes in the traders' market selection strategies, and in order to remain competitive, eventually, they have to charge the lowest profit fee. However, in the static analysis, the competing marketplaces set their fees once at the beginning, and their charging strategies do not respond to the changes in the traders' market selection strategies.

In the previous analysis (see Section 3.2.3.2), we introduced randomisation for the traders' market selection strategies to analyse the effect of exploration and bounded rationality. Now we do the same for the above co-evolutionary setting. In doing so, we find that, as the probability of

[17]If a lower minimal profit fee would have been allowed, e.g. 1%, then both competing marketplaces will converge to this lower profit fee.

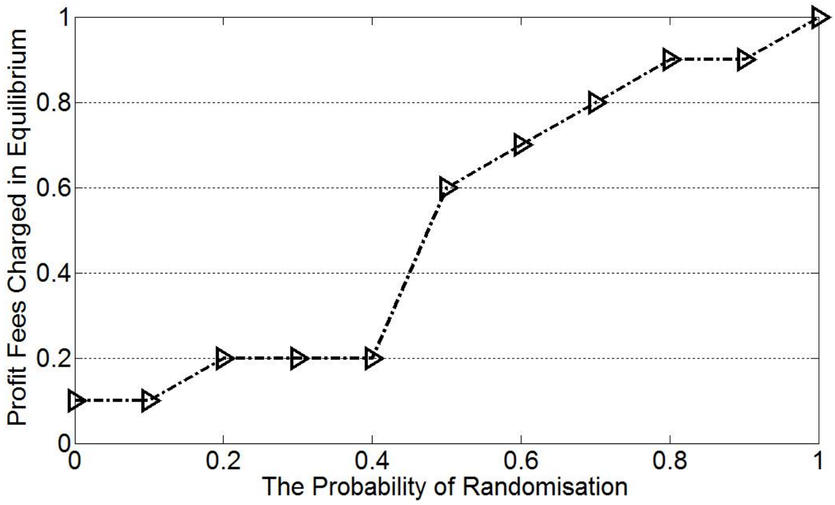

FIGURE 3.18: Equilibrium charging strategies of marketplace 1 and 2 with respect to randomisation of market selection when two competing marketplaces have identical initial charging strategies.

traders randomly choosing marketplaces increases, in equilibrium, fees increase. For example, when we introduce 20% randomisation, then in equilibrium, both marketplaces will charge a 20% profit fee. The result is shown in Figure 3.18. From this, we can see when randomisation goes above 50% (i.e. when traders have a very high probability of randomising their market selection), in equilibrium, the marketplaces charge very high fees. Especially, when the randomisation reaches 100%, both marketplaces charge 100% profit fee. This is because two identical competing marketplaces have the same evolutionary process, and thus they cannot attract traders from each other. When traders have probabilities of randomly choosing marketplace, both marketplaces will find that, even though they charge higher fees, they still keep traders. Thus marketplaces will charge higher fees to make more profits.

Two marketplaces initially having different charging strategies:

Now we consider the more general case, in which the marketplaces have different initial charging strategies. For example, we assume that initially marketplace 2 is slightly cheaper than marketplace 1, and thus marketplace 1 is slightly disadvantaged in terms of the traders' probability of choosing marketplaces. For this setting, the evolutionary charging strategies of marketplace 1 and 2 are shown by Figures 3.19(a) and 3.19(b), and the dynamic changes of the traders' probabilities of choosing marketplace 1 are shown in Figure 3.19(c). From these, we can see that in equilibrium, all traders converge to marketplace 1, which is initially disadvantaged. The reason is that from Figures 3.19(a) and 3.19(b), we can see that marketplace 1 decreases its fees to attract traders because of its disadvantageous position in the initial state, and marketplace 2 increases fees since it has an advantageous position in the initial state. Although there exist small fluctuations for traders' probabilities of choosing marketplaces because of the fee changes of marketplace 1 and 2, eventually all traders will converge to marketplace 1. This shows that it is possible for an initial disadvantaged marketplace to beat an advantaged one by dynamically adapting its fees. We also find that once marketplace 1 attracts all traders, it will charge higher

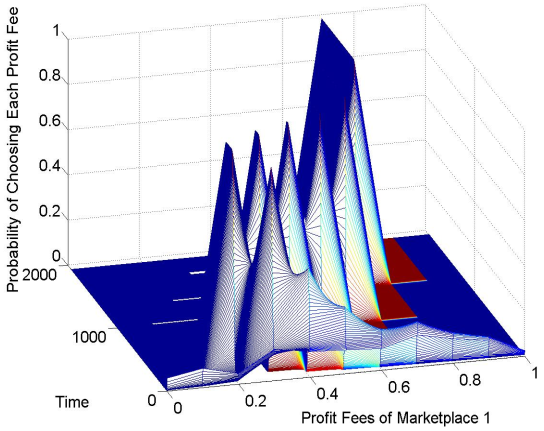

(a) Charging strategy of marketplace 1.

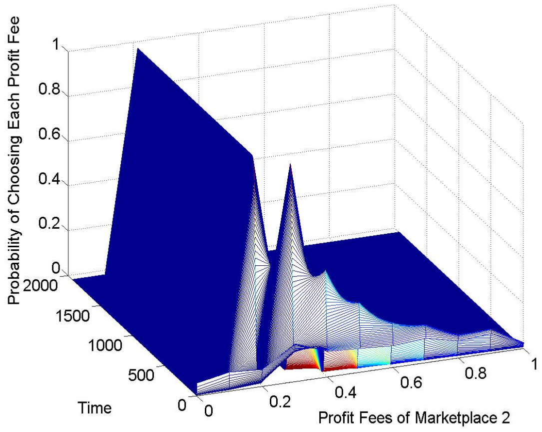

(b) Charging strategy of marketplace 2.

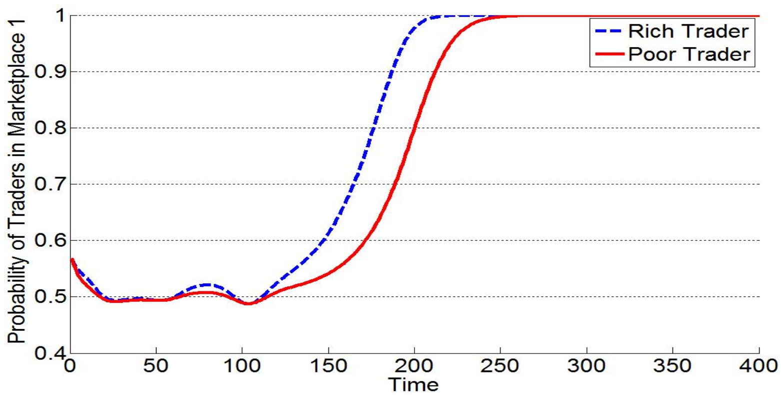

(c) Traders' probability of choosing marketplace 1.

FIGURE 3.19: Evolutionary process of charging strategies of marketplace 1 and 2 and traders' probability of choosing marketplace 1 when two competing marketplaces have different initial charging strategies.

fees, around 70% profit fee, but still keep traders. This is because the profit fees of both competing marketplaces are within the lock-in region, and therefore marketplace 1 can keep traders even though it is more expensive. However, if we again introduce randomisation (see Figure 3.20), we see that the behaviour of the traders' market selection changes significantly. In detail, we can see that marketplace 1 tries to charge a higher profit fee, but because of random exploration, traders will migrate to marketplace 2 (the cheaper marketplace). This causes marketplace 1 to reduce its fees and traders to migrate back to this marketplace. In fact, we observe that the strategies of the traders and the marketplaces never converge to an equilibrium. However, by observing the overall evolutionary process, we still can see that, on average, marketplace 1 charges slightly higher fees than marketplace 2 because of its higher initial market share.

Two different marketplaces having an adaptive charging strategy and a fixed charging strategy respectively:

In the above, we analysed the dynamic behaviour of both competing marketplaces evolving their

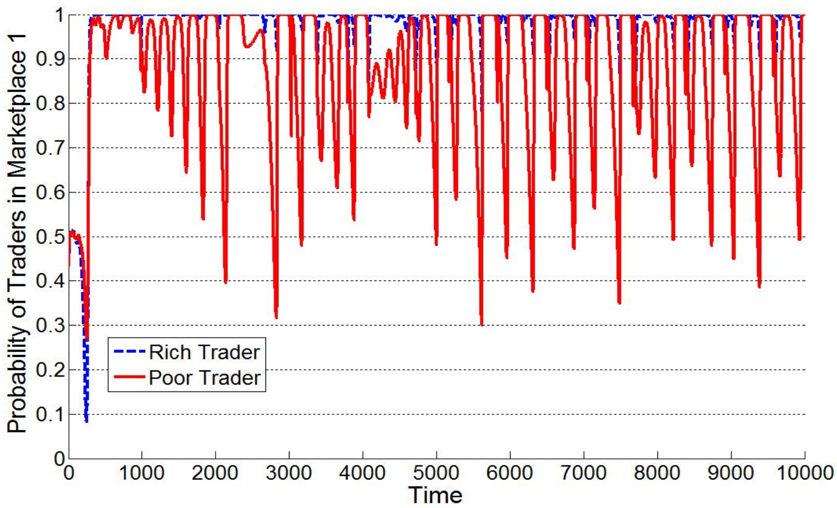

FIGURE 3.20: Traders' probability of choosing marketplace 1 with 20% random exploration when two competing marketplaces have different initial charging strategies .

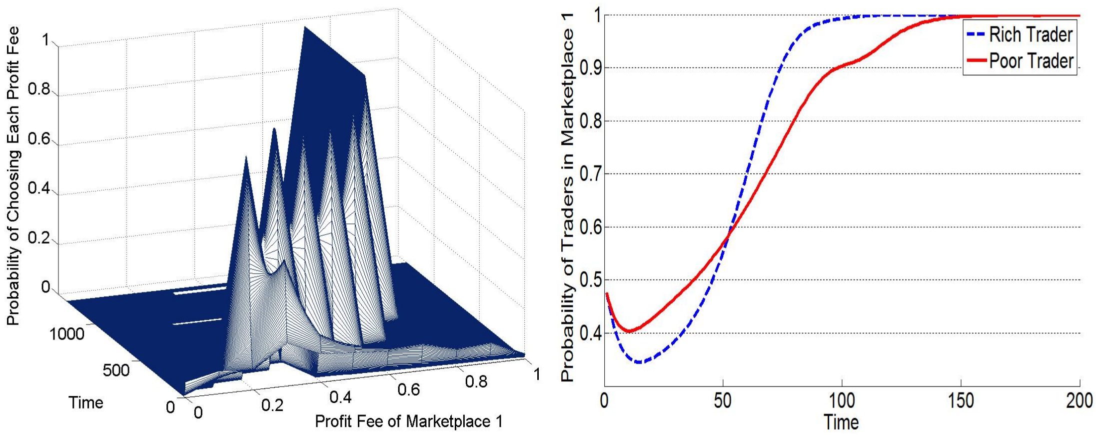

(a) Charging strategy of marketplace 1.

(b) Traders' probability of choosing marketplace 1.

FIGURE 3.21: Evolutionary process of charging strategy of marketplace 1 and traders' probability of choosing marketplace 1 when marketplace 1 adopts an adaptive charging strategy and marketplace 2 adopts a fixed charging strategy.

charging strategies. In the real world, however, some marketplaces may adopt a fixed charging strategy, which means that they will not change fees during a specific time. To consider this situation, we now analyse how a marketplace with an adaptive charging strategy competes with a marketplace with a fixed charging strategy. As an example, we assume that marketplace 2 fixes its profit fee at 30%, and marketplace 1 evolves its charging strategy and initially marketplace 1 is slightly more expensive than marketplace 2. The evolutionary process of the charging strategy of marketplace 1 is shown in Figure 3.21(a), and the dynamic changes of the traders' probability of choosing marketplace 1 for this setting are shown in Figure 3.21(b). From these figures, we can see that initially, marketplace 1 decreases its fee, and when attracting all traders, it will increase its fee, but still keep traders. In equilibrium, marketplace 1 will charge 70% profit fee. However, if in the beginning marketplace 1 is much more expensive than marketplace

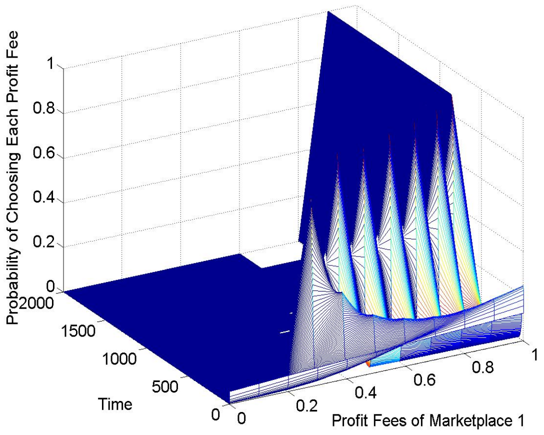

(a) Charging strategy of marketplace 1.

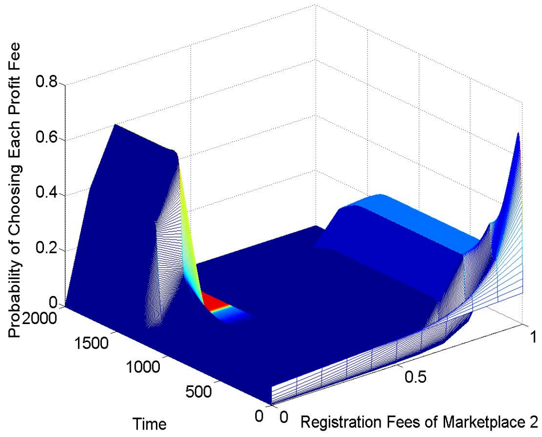

(b) Charging strategy of marketplace 2.

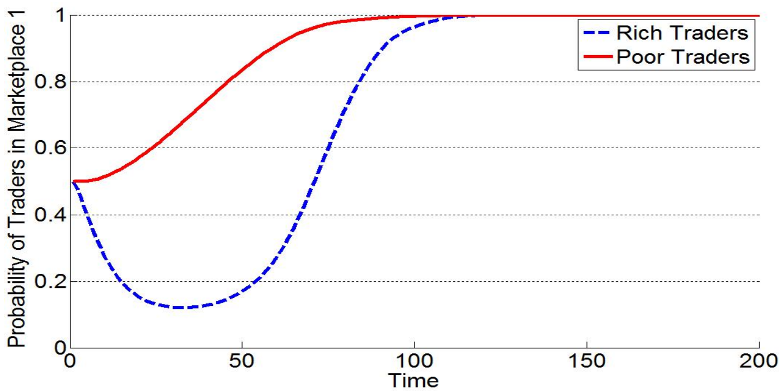

(c) Traders' probability of choosing marketplace 1.

FIGURE 3.22: Evolutionary process of charging strategies of marketplace 1 and 2 and traders' probability of choosing marketplace 1 when marketplace 1 charges a profit fee and marketplace 2 charges a registration fee.

2 (which means that marketplace 2 has a very large lock-in region), then even though marketplace 1 charges a very low fee, it still fails to attract traders. If both marketplaces have similar initial charging strategies, then the marketplace using an adaptive charging strategy can beat the marketplace using a fixed charging strategy. However, when a marketplace initially has a large market share, it is difficult for a new marketplace to obtain market share, even when undercutting its competitors.

Two different marketplaces charging different types of fees:

Finally we analyse how marketplaces evolve their charging strategies when different types of fees are charged. Previously (in Section 3.2.3.2), when marketplaces charge different types of fees, we showed that rich traders prefer marketplaces that charge registration fees, and poor traders prefer marketplaces that charge profit fees. From Figure 3.22(c), we find that, initially, rich traders still prefer marketplace 2 charging a registration fee, and poor traders prefer marketplace 1 charging a profit fee, which is the same as the previous analysis. However, from Figures

3.22(b) and 3.22(c), we can see that when rich traders choose marketplace 2, this marketplace charges a higher registration fee. Then, in contrast to the previous analysis, where different types of traders converge to different marketplaces in equilibrium, rich traders leave marketplace 2, and all traders converge to marketplace 1. Once all traders choose marketplace 1, from Figure 3.22(a), we can see marketplace 1 charges 90% profit fee, but still keeps the traders.

3.4 Summary

In this chapter, we proposed a game-theoretic framework for analysing competing double auction marketplaces that vie for traders and make profits by charging fees. *Firstly*, we analysed the equilibrium market selection strategies for traders for a given fee system. In more detail, we used game theory to analyse the equilibrium market selection strategies and adopted evolutionary game theory to investigate how traders dynamically change their strategies, and thus, which equilibrium, if any, can be reached. In so doing, we showed that when the same type of fees are charged by two marketplaces, it is unlikely that competing marketplaces will continue to co-exist when traders converge to their equilibrium market selection strategies. Eventually, all the traders will congregate in one marketplace. However, when different types of fees are allowed (registration fees and profit fees), competing marketplaces are more likely to co-exist in equilibrium, where rich traders will converge to the marketplace charging a registration fee, and poor traders will converge to the marketplace charging a profit fee. Somewhat surprisingly, we found that sometimes all the traders eventually migrate to the marketplace that charges higher fees. Thus we further analysed this phenomenon, and specifically analysed how random exploration by traders affects this migration.

Secondly, we analysed the equilibrium charging strategies of the marketplaces using two different approaches. In the first, we derived the equilibrium charging strategies by a static analysis. However, this approach did not consider the interaction between traders' strategies and marketplaces' strategies. We tackled this limitation by using a co-evolutionary approach to analyse this game. Specifically, we considered the competition of the marketplaces as a two-stage game, where the traders' market selection strategies and marketplaces' charging strategies affect each other. In particular, we used a co-evolutionary approach to analyse how competing marketplaces dynamically set fees while taking into account the dynamics of the traders' market selection strategies. In so doing, we found that two initially identical marketplaces undercut each other, and they will eventually charge the minimal fee that guarantees positive market profits for them. Furthermore, we also extended the co-evolutionary analysis of the marketplaces' charging strategies to more general cases. We found that by dynamically evolving the charging strategy, it is possible for the marketplace that is initially at a disadvantage to outperform its opponent, which is also able to evolve its charging strategy. Furthermore, we showed that an initially disadvantaged marketplace with an adaptive charging strategy can beat the initially advantaged one with a fixed charging strategy.

Furthermore, we also want to point out that a certain conclusion drawn from this chapter is consistent with the work done by Sohn et al. (2010); Niu et al. (2008a) in the context of the CAT competition. In more detail, both of them presented that intra-marginal traders will select the marketplace charging a registration fee with a high probability, and extra-marginal traders will select the marketplace charging a profit fee with a high probability. Their conclusion aligns well with our conclusion that when two double auction marketplaces charge different types of fees, it can happen that rich traders (corresponding to intra-marginal traders) converge to the marketplace charging a registration fee, and poor traders (corresponding to extra-marginal traders) converge to the marketplace charging a profit fee.

The work in this chapter addresses our research challenges of analysing equilibrium market selection strategies for traders and equilibrium charging strategies for marketplaces (i.e. research challenges 1 and 4, see Section 1.2), which is the *first* theoretical work in the context of competing double auction marketplaces. This analysis is insightful and can be used to guide the design of a charging strategy. For example, the lock-in region gives us the insight that the competing marketplace should initially charge lower and even no fees to attract traders, and once it has built up a larger market share, the marketplace can charge fees to make profits, but still can maintain market share at a good level. However, in this work, we only considered discrete trader types, and assumed that traders adopt a simple, truth-telling bidding strategy. In the next chapter, we will address these shortcomings by considering continuous trader types and analysing both equilibrium market selection and bidding strategies for traders. We will also extend this analysis to the settings with multi-home trading and hybrid trading, and with different properties of goods.

Chapter 4

Analysis of Competing Marketplaces with Continuous Trader Types

In Chapter 3, we game theoretically analysed the equilibrium market selection strategies for traders and equilibrium charging strategies for marketplaces. Now, while this analysis led to a number of important insights, it was restricted to the setting with two discrete trader types. Furthermore, we made the simplifying assumption that traders use a truth-telling bidding strategy, which means that traders will submit their types as their shouts. However, in practice, trader types are often drawn from a wider range, and they may bid strategically (such as shading shouts in order to make more profits). In addition, we assumed that marketplaces can only choose from two types of fees (registration and profit) to make profit. However, as we will show in this chapter, when traders can bid strategically, we find that neither fees are effective in terms of making profits or keeping traders (see Section 4.4). Charging registration fees will cause traders to leave the marketplace quickly, and charging profit fees encourages traders to hide their actual profits by shading their shouts, resulting in very low revenue for the marketplaces. In this chapter, we will address all these limitations. Specifically, we consider a setting with continuous trader types, where we use fictitious play (a computational learning approach) to analyse both the equilibrium market selection and bidding strategies for traders. Furthermore, because of ineffectiveness of both registration and profit fees in making profits and keeping traders, we consider two more types of fees: transaction and transaction price percentage fees. We introduce the former one because it is used in the CAT competition (see Section 2.5.1), and we introduce the latter one because it is commonly used in real-world auctions (e.g. both eBay and Amazon charge a percentage fee on the final sale prices to sellers). In so doing, we analyse what types of fees are effective in making profits and keeping traders, and analyse how competing marketplaces set fees in equilibrium.

We furthermore extend the model in several directions. First, in addition to the single-home trading environment where traders can only enter one marketplace at a time (the analysis in Chapter 3 was restricted to this environment), we also consider other trading environments: multi-home environments where traders can enter multiple marketplaces at a time, and hybrid environments where one side of traders (e.g. buyers) use multi-home trading and the other side of traders (e.g. sellers) use single-home trading. These different trading environments affect the way in which traders select marketplaces and submit shouts. For example, in the multi-home trading environment, traders will enter multiple marketplaces if they can make positive profits in these marketplaces. Furthermore, if buyers can enter multiple marketplaces at a time, they can obtain multiple goods. As discussed in Section 1.2, these goods can be either independent, substitutes, or complementary. These different properties also affect traders' strategies. For example, when trading complementary goods, buyers will prefer to buy as many goods as they can, and thus they will try to bid high in several marketplaces to make more transactions. Therefore, in this chapter, in addition to considering different trading environments, we will also analyse goods with these different properties.

Specifically, this chapter addresses research challenges 1 (analysing market selection strategies for traders), 2 (analysing bidding strategies for traders) and 4 (analysing charging strategies for marketplaces). The structure of this chapter is as follows. In Section 4.1, we formalise the above setting, and derive expected utilities of traders and marketplaces in this new setting. In Section 4.2, we describe the fictitious play (FP) algorithm used in our analysis. In Section 4.3, we use FP to analyse traders' equilibrium market selection and bidding strategies in different trading environments with different good properties. In Section 4.4, we analyse the effects of different types of fees on obtaining market profits and how competing marketplaces should set fees in equilibrium. Finally, we summarise in Section 4.5

4.1 Framework

In this section, we first extend the setting from Chapter 3, and then derive equations to calculate expected utilities of traders and marketplaces for this new setting.

4.1.1 Basic Settings

The basic setting for traders and marketplaces in this chapter is similar to that used in Chapter 3 (see Section 3.1.1). In this section we will discuss the changes with respect to Chapter 3 as a result of the changes in the setting described above. First of all, since we consider continuous trader types, in this setting we assume that types of buyers (sellers) are identically and independently drawn from the cumulative distribution function F^b (F^s) with support [0,1] (in contrast to that in Chapter 3 where we only consider discrete trader types). Furthermore, we consider two more types of fees: transaction and transaction price percentage. Consequently, the fee structure of a marketplace m is now defined as $p_m = (r_m, t_m, q_m, o_m)$, $r_m \geq 0$, $t_m \geq 0$, $q_m \in [0, 1]$ and $o_m \in [0, 1]$, where r_m is a registration fee charged to traders when they enter the marketplace, t_m is a transaction fee charged to buyers and sellers when they make transactions, q_m is a profit fee charged on profits made by buyers and sellers, and o_m is a transaction price percentage fee charged on the transaction price of buyers and sellers. As before, the fee structures of all competing marketplaces constitute a fee system $\bar{P} = \langle p_1, p_2, ...p_M \rangle$. Moreover, we further make an assumption that traders will incur a small cost ε when they choose any marketplace (such as time cost for trading online, travel and time costs for trading in shopping mall). We do this so that they slightly prefer choosing no marketplace than choosing a marketplace and making no transactions (even if $r_m = 0$). This small cost will help us distinguish buyers' behaviour between bidding zero and not choosing the marketplace, and sellers' behaviour between bidding one and not choosing the marketplace.

The process of a trading round is the same as Chapter 3 (see Section 3.1) except that, in this setting, traders can enter multiple marketplaces when multi-home trading is allowed. As before, we assume that each trader can only trade a single unit of the good in each marketplace. However, because now multi-home trading is allowed, traders may trade multiple goods when they enter multiple marketplaces. We assume that goods traded in different marketplaces are identical. As we said before, these goods can be either independent, substitutes, or complementary. We model these different preferences as follows. For a buyer with type θ^b, the value that it derives when it successfully purchases T units of goods is given by:

$$v^b(\theta^b, T) = \alpha_T^b * \theta^b \tag{4.1}$$

where we refer to α_T^b as the "buyer preference coefficient", determining whether the buyers have independent, substitutable or complementary preferences. In more detail, if goods are independent for the buyer, then the values for individual good are additive, i.e. $\alpha_1^b = 1$ and $\alpha_T^b - \alpha_{T-1}^b = 1, \forall T \geq 2$ (which is equivalent to $\alpha_T^b = T$). If, on the other hand, the goods are substitutes, then the total value for getting T goods is subadditive, i.e. $\alpha_1^b = 1$ and $0 \leq \alpha_T^b - \alpha_{T-1}^b < 1, \forall T \geq 2$. In particular, for perfectly substitutable goods for the buyer, we have $\alpha_T^b = 1$ $(T = 1, ..., M)$. Finally, if goods are complementary for the buyer, the total value is superadditive, i.e. $\alpha_1^b = 1$ and $\alpha_T^b - \alpha_{T-1}^b > 1, \forall T \geq 2$. Specifically, for perfectly complementary goods, we have $\alpha_T^b = 0$ $(T = 1, ..., M-1)$ and $\alpha_M^b = 1$, i.e. the buyer obtains value θ^b when it successfully purchases M goods, and the buyer obtains zero value when it purchases less than M goods. The value function for a seller is defined analogously, where α_T^s is the seller preference coefficient. We assume that α_T^b and α_T^s are the same for all buyers and sellers, and these parameters are common knowledge.

In addition to multi-home, we now allow traders to bid strategically. In what follows, we describe how this affects the framework from Chapter 3. Specifically, we make the assumption that there is a finite number of bids and asks and that these are discrete. The reason for doing so is two-fold. First of all, 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. For convenience, we further assume that buyers and sellers have the same shout space, which is given by $\Phi = \{0, \frac{1}{D}, \frac{2}{D}, ..., \frac{D-1}{D}, 1\} \cup \{\ominus\}$, i.e. the bid(ask) space comprises $D + 1$ allowable bids(asks) from 0 to 1 with step size $1/D$, and $\ominus$ means not submitting a shout in the marketplace (i.e. not choosing the marketplace). Recall that, in our new setting, traders can place shouts in multiple marketplaces. We refer to a combinational shout across multiple marketplaces as an *action*. Formally, a buyer's action is defined as a tuple $\delta^b = \langle d_1^b, d_2^b, ..., d_M^b \rangle$, where the buyer bids d_m^b in marketplace m if $d_m^b \neq \ominus$, and does not choose marketplace m if $d_m^b = \ominus$. Similarly, a seller's action is given by $\delta^s = \langle d_1^s, d_2^s, ..., d_M^s \rangle$. The set of all possible actions constitutes the *action space*, which is defined as $\Delta = \Phi^M$. Note that in our system, both buyers and sellers have the same action space.

Now, a trader's action is determined by its strategy, which is a mapping from the set of types to the action space. In addition, each trader does not know what the exact types of the other traders are, and only know their type distribution functions. Therefore, given the trader's strategy and the type distribution function, we can derive the probability of a certain action being played by this trader. Since the expected utility of a trader (and a marketplace) is directly dependent on its beliefs about other traders' action choices, instead of looking at traders' strategies, in what follows we directly consider traders' action distributions. This will also be convenient when we use fictitious play to derive traders' equilibrium strategies (see Section 4.2). Specifically, we change our notation from before and now use ω_i^b (ω_i^s) to denote the probability of action δ_i^b (δ_i^s) being chosen by a buyer (seller). Furthermore, we use $\Omega^b = (\omega_1^b, \omega_2^b, ..., \omega_{|\Delta|}^b)$, $\sum_{i=1}^{|\Delta|} \omega_i^b = 1$, to represent the probability distribution of buyers' actions, and $\Omega^s = (\omega_1^s, \omega_2^s, ..., \omega_{|\Delta|}^s)$ for the sellers' action distribution.

4.1.2 The Trader's Expected Utility

Before analysing the strategies of the traders and the marketplaces, we first need to derive equations to calculate expected utilities of traders and marketplaces. Although in Chapter 3 we have given equations to calculate them, these equations are based on the assumption that traders submit their types as shouts, and thus are not appropriate in this setting. Therefore, in this section, we derive the equations to calculate the expected utility of a trader based on the action distribution of other traders. Then in the following section, we will describe how to calculate the expected utilities of 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 action choices of other traders. In the following, we calculate the expected utility of a buyer with type θ^b adopting the action $\delta^b = \langle d_1^b, d_2^b, ..., d_M^b \rangle$ given the other buyers' action distribution Ω^b and the sellers' action distribution Ω^s, and the fee system $\bar{P}$. The expected utility consists of two parts: the *expected value* on the goods, and the *expected payment*. In the following, we derive these two parts respectively.

To calculate the buyer's expected value, since we consider the equilibrium matching policy matching buyers having high bids with sellers having low asks, we need to consider the buyer's positions in the available marketplaces (the buyer's position is the rank of the buyer's bid among all descendingly sorted bids). Clearly, a buyer's position 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 the probabilistic information about the buyer's position in a marketplace. Furthermore, since buyers can place multiple bids in multiple marketplaces at the same time, bids in different marketplaces are correlated with each other, and thus the buyer's positions in different marketplaces are also correlated with each other. Therefore, we cannot consider each marketplace independently and need to consider the buyer's *joint positions* in marketplaces. Furthermore, since we consider discrete bids and multiple buyers may place the same bids, we need to use a tie-breaking rule to determine a buyer's position.

Therefore, to calculate the expected value, we need to take the following steps. First, we calculate the expected joint positions of the buyer by considering all possible action choices of other buyers and their action distributions. In addition to other buyers' actions, a buyer's expected value also depends on the sellers' actions. Therefore, we then consider the number of sellers choosing different actions. Finally, given the buyer's joint positions and the number of sellers choosing different actions, we can derive the buyer's expected value by considering all possible numbers of units it can win. In the following, we discuss these steps in turn.

Firstly, we describe how to determine a buyer's joint positions across the marketplaces. As we said above, when we know the number of buyers choosing different actions, we can determine the buyer's joint positions. Specifically, we use a $|\Delta|$-tuple $\bar{x} = \langle x_1, ...x_{|\Delta|}\rangle \in \mathcal{X}$ to represent the number of buyers choosing different 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:

$$\rho^b(\bar{x}) = \binom{B-1}{x_1, ..., x_{|\Delta|}} * \prod_{i=1}^{|\Delta|} \left(\omega_i^b\right)^{x_i} \tag{4.2}$$

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

$$X_m^{>}(\bar{x}, d_m^b) = \sum_{\delta_i^b \in \Delta : d_{im}^b > d_m^b} x_i \tag{4.3}$$

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):

$$X_m^{=}(\bar{x}, d_m^b) = \sum_{\delta_i^b \in \Delta : d_{im}^b = d_m^b} x_i \tag{4.4}$$

Due to having 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 $X_m^>(\bar{x}, d_m^b)+X_m^=(\bar{x}, d_m^b)+1$, which constitutes the buyer's position range in this marketplace. Since $X_m^=(\bar{x}, d_m^b)+1$ buyers have the same bid, as we said previously, a tie-breaking rule is needed to determine the buyer's position. Here we adopt a standard rule where each of these possible positions[1] 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 a M-tuple $\bar{v}_{\bar{x}} = \langle v_1, ..., v_M\rangle \in \mathcal{V}_{\bar{x}}$ to represent one of the possible joint positions where 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 $X_m^>(\bar{x}, d_m^b)+1 \le v_m \le X_m^>(\bar{x}, d_m^b)+X_m^=(\bar{x}, d_m^b)+1$ $(m = 1, ..., M)$. The probability of the buyer having the joint positions $\bar{v}_{\bar{x}}$ given the tuple $\bar{x}$ is:

$$\Phi(\bar{v}_{\bar{x}}) = \prod_{m=1}^{M} \frac{1}{X_m^=(\bar{x}, d_m^b)+1} \tag{4.5}$$

Note that tie-breaking occurs independently for each marketplace.

In addition to depending on positions in different marketplaces, the buyer's expected value also depends on sellers' action choices. Specifically, we use a $|\Delta|$-tuple $\bar{y} = \langle y_1, ...y_{|\Delta|}\rangle \in \mathcal{Y}$ 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:

$$\rho^s(\bar{y}) = \binom{S}{y_1, ..., y_{|\Delta|}} * \prod_{i=1}^{|\Delta|} \left(\omega_i^s\right)^{y_i} \tag{4.6}$$

Now 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 value on traded goods. Since the buyer can enter multiple marketplaces and thus purchase multiple goods, we need to consider its expected value on different units of goods. Remember that 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 4.1.1, we have defined the buyer's value $v^b(\theta^b, T)$ on T units of goods by considering different good properties (see Equation 4.1). Now by considering all possible marketplace subsets with cardinality T, where exactly T transactions are made by this buyer, we obtain the buyer's expected value when it purchases T units of goods given its joint positions $\bar{v}_{\bar{x}}$ and the number of sellers choosing different actions $\bar{y}$:

$$\begin{aligned}\tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) &= \sum_{\mathcal{M}_I \subset 2^{\mathcal{M}}:|\mathcal{M}_I|=T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) * v^b(\theta^b, T) \\ &= \sum_{\mathcal{M}_I \subset 2^{\mathcal{M}}:|\mathcal{M}_I|=T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) * \alpha_T^b * \theta^b \end{aligned} \tag{4.7}$$

where $\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I)$ indicates whether the buyer makes transactions in marketplaces $\mathcal{M}_I$ and does not make transactions in $\mathcal{M} - \mathcal{M}_I$. Note that conditional on the buyer's joint positions

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

and the number of sellers choosing different actions, whether the buyer making a transaction or not in each marketplace is independent of each other, and thus whether the buyer making transactions in $\mathcal{M}_I$ and not making transactions in $\mathcal{M} - \mathcal{M}_I$ is given by:

$$\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) = \prod_{m \in \mathcal{M}_I} \psi^b(v_m, \bar{y}, m, d_m^b) * \prod_{m \in \mathcal{M} - \mathcal{M}_I} \chi^b(v_m, \bar{y}, m, d_m^b) \tag{4.8}$$

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 make a transaction in marketplace m. Given the number of sellers choosing different actions, we will know what asks are placed in marketplace m, from which we can calculate the number of asks which are not greater than d_m^b:

$$Y_m^{\leq}(\bar{y}, d_m^b) = \sum_{\delta_i^s \in \Delta : d_{im}^s \leq d_m^b} y_i \tag{4.9}$$

Now whether the buyer making a transaction in marketplace m is given by:

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

and whether the buyer not making a transaction in marketplace m is:

$$\chi^b(v_m, \bar{y}, m, d_m^b) = \begin{cases} 1 & \textbf{\textit{if }} Y_m^{\leq}(\bar{y}, d_m^b) < v_m \\ 0 & \textbf{\textit{if }} Y_m^{\leq}(\bar{y}, d_m^b) \geq v_m \end{cases}$$

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 value is given by:

$$\tilde{V}(\theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{\bar{x} \in \mathcal{X}} \rho^b(\bar{x}) * \sum_{\bar{v}_{\bar{x}} \in \mathcal{V}_{\bar{x}}} \Phi(\bar{v}_{\bar{x}}) * \sum_{\bar{y} \in \mathcal{Y}} \rho^s(\bar{y}) * \sum_{T=1}^{M} \tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) \tag{4.10}$$

After deriving the expected value, in the following, we derive the expected payment of the buyer given the action distributions of buyers and sellers, Ω^b and Ω^s, and the fee system $\bar{P}$. Firstly, we derive the buyer's expected payment given its joint positions $\bar{v}_{\bar{x}}$ and the number of sellers choosing different actions $\bar{y}$. This is equal to the sum of the expected payment in each marketplace, which is:

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

where $\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, p_m)$ 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 the bid d_m^b. We denote the ask matched with bid d_m^b in marketplace m as d_m^s. As a result, $\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, p_m)$ is given by:

$$\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, p_m) = \begin{cases} 0 & \textit{if } d_m^b = \ominus \\ \tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, d_m^s, p_m) + r_m + \varepsilon & \textit{if } d_m^b \geq d_m^s \\ r_m + \varepsilon & \textit{otherwise} \end{cases}$$

where $\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, d_m^s, p_m)$ is the buyer's expected payment in marketplace m excluding registration fee r_m and constant cost ε when it is matched with d_m^s, which is given by:

$$\tilde{\mathcal{P}}_m^b(v_m, \bar{y}, \theta^b, d_m^b, \Omega^b, \Omega^s, d_m^s, p_m) = \text{TP} + t_m + \text{TP} * o_m + (d_m^b - \text{TP}) * q_m \tag{4.11}$$

where $\text{TP} = d_m^s * k_m + d_m^b * (1 - k_m)$ is the transaction price, t_m is the transaction fee, $\text{TP} * o_m$ is the payment of transaction price percentage fee, and $(d_m^b - \text{TP}) * q_m$ is the payment of profit fee.

Now by considering 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 payment is given by:

$$\tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) = \sum_{\bar{x} \in \mathcal{X}} \rho^b(\bar{x}) * \sum_{\bar{v}_{\bar{x}} \in \mathcal{V}_{\bar{x}}} \Phi(\bar{v}_{\bar{x}}) * \sum_{\bar{y} \in \mathcal{Y}} \rho^s(\bar{y}) * \tilde{\mathcal{P}}^b(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) \tag{4.12}$$

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

$$\tilde{U}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) = \tilde{V}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) - \tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) \tag{4.13}$$

Now we have derived equations to calculate expected utilities of traders. This way of calculating traders' expected utilities is intuitive and thus can be easily explained. However, the computation is heavy since we need to consider all the possible numbers of traders choosing different actions. In more detail, given shout space Φ and M competing marketplaces, there are $|\Phi|^M$ actions for traders. Then for B buyers, there are $|\mathcal{X}| = |\Phi|^{M*(B-1)}$ possibilities for buyers' action choices (excluding the buyer itself), and $|\mathcal{Y}| = |\Phi|^{M*S}$ for sellers. In Appendix B, in order to reduce the computation, we introduce an alternative approach to calculate traders' expected utilities (which is somewhat less intuitive and thus is relatively harder to explain). Specifically, we compare buyers' actions in terms of comparing bids in these actions. By doing so, we obtain 3^M different events for buyers' action comparison. By considering the number of buyers choosing actions satisfying each event, we can calculate the buyer's joint positions. Now we reduce the possibilities of buyers' action choices from $|\Phi|^{M*(B-1)}$ to $3^{M*(B-1)}$. For sellers, instead of directly looking at all possible numbers of sellers choosing different actions, we consider the possibilities of the number of the sellers in different marketplaces. By doing so, we reduce the possibilities from $|\Phi|^{M*S}$ to $(S+1)^M$ (note that the number of sellers in each marketplace is from 0 to S since multi-home trading is allowed)[2]. For details of this approach, we refer to Appendix B. Note

[2]For example, in the following analysis, we consider 11 possible shouts plus $\ominus$, 2 competing marketplaces, 5

that in our implementation of the algorithm, in order to reduce the computation, we adopt the approach in Appendix B to calculate a trader's expected utility. However, both approaches are mathematically equivalent.

4.1.3 The Marketplace's Expected Utility

After deriving equations to calculate the expected utilities of traders, we now calculate expected utilities of marketplaces. Specifically, in the following, we derive equations to calculate marketplace m's expected utility given its fee structure p_m and the action distributions of buyers and sellers, Ω^b and Ω^s. Intuitively, we can see that the expected utility depends on the number of traders choosing each action. Similarly, we use a $|\Delta|$-tuple $\bar{x} = \langle x_1, ..., x_{|\Delta|}\rangle \in \mathcal{X}'$, $\sum_{i=1}^{|\Delta|} x_i = B$, to denote the number of buyers choosing different actions, where x_i is the exact number of buyers choosing action δ_i^b, and $\mathcal{X}'$ is the set of all such possible tuples. We use $\bar{y} = \langle y_1, ..., y_{|\Delta|}\rangle \in \mathcal{Y}$, $\sum_{i=1}^{|\Delta|} y_i = S$, to denote the number of sellers choosing different actions. Given the number of buyers and sellers choosing different actions, $\bar{x}$ and $\bar{y}$, we will know the bids and asks placed in marketplace m. Then marketplace m's expected utility is calculated as follows. Since marketplace m uses equilibrium matching to match traders, we first sort the bids descendingly and asks ascendingly in marketplace m, and then match high bids with low asks. Specifically, we assume that there are T transactions in total in marketplace m, and in transaction t, we use TP_t, Λ_t^b and Λ_t^s to represent the transaction price, the buyer's share of the trading surplus[3], and the seller's share of the trading surplus respectively. These can be easily calculated. For example, for a transaction made by buyer with bid d_t^b and seller with ask d_t^s, the transaction price is $\mathrm{TP}_t = d_t^s * k_m + d_t^b * (1-k_m)$, the buyer's share of trading surplus is $\Lambda_t^b = d_t^b - \mathrm{TP}_t = (d_t^b - d_t^s) * k_m$, and the seller's share of trading surplus is $\Lambda_t^s = \mathrm{TP}_t - d_t^s = (d_t^b - d_t^s) * (1-k_m)$. The marketplace's utility is:

$$U_m(p_m, \bar{x}, \bar{y}) = \sum_{\delta_i^b \in \Delta : d_{im}^b \neq \ominus} x_i * r_m + \sum_{\delta_i^s \in \Delta : d_{im}^s \neq \ominus} y_i * r_m + \sum_{t=1}^{T} (t_m * 2 + \Lambda_t^b * q_m + \Lambda_t^s * q_m + \mathrm{TP}_t * o_m * 2) \quad (4.14)$$

where d_{im}^b is the bid placed in marketplace m through the action δ_i^b. In this equation, the former two parts are profits from charging registration fees to buyers and sellers respectively, and the last part is the profit from charging transaction fees, profit fees and transaction price percentage fees.

Now we have obtained the marketplace's expected utility given the number of buyers and sellers choosing different actions, $\bar{x}$ and $\bar{y}$. Furthermore, the probability of $\bar{x}$ appearing is:

$$\varrho_m^b(\bar{x}) = \binom{B}{x_1, ..., x_{|\Delta|}} * \prod_{i=1}^{|\Delta|} \left(\omega_i^b\right)^{x_i} \quad (4.15)$$

buyers and 5 sellers. By using this alternative approach, we reduce the possibilities of buyers' action choices from $12^8 = 429981696$ to $3^8 = 6561$, and reduce the possibilities for sellers from $12^{10} = 61917364224$ to $6^2 = 36$.

[3]The trading surplus of a transaction is the difference of the matched bid and ask in this transaction.

and the probability of $\bar{y}$ appearing is:

$$\varrho_m^s(\bar{y}) = \binom{S}{y_1, ..., y_{|\Delta|}} * \prod_{i=1}^{|\Delta|} \left(\omega_i^s\right)^{y_i} \tag{4.16}$$

At this moment, we can get the marketplace's expected utility given action distributions Ω^b and Ω^s:

$$\tilde{U}_m(p_m, \Omega^b, \Omega^s) = \sum_{\bar{x}\in\mathcal{X}'} \varrho_m^b(\bar{x}) * \sum_{\bar{y}\in\mathcal{Y}} \varrho_m^s(\bar{y}) * U_m(p_m, \bar{x}, \bar{y}) \tag{4.17}$$

Furthermore, given action distributions Ω^b and Ω^s, we can obtain the expected number of traders choosing marketplace m, which we will use when analysing the effects of different types of fees on making profits and keeping traders (see Section 4.4.1). Specifically, given the number of buyers and sellers choosing different actions, $\bar{x}$ and $\bar{y}$, the number of traders in marketplace m is:

$$Q_m(\bar{x}, \bar{y}) = \sum_{\delta_i^b\in\Delta:d_{im}^b\neq\ominus} x_i + \sum_{\delta_i^s\in\Delta:d_{im}^s\neq\ominus} y_i \tag{4.18}$$

Then by considering all possible numbers of buyers and sellers choosing different actions, we obtain the expected number of traders in marketplace m:

$$\tilde{Q}_m(\Omega^b, \Omega^s) = \sum_{\bar{x}\in\mathcal{X}'} \varrho_m^b(\bar{x}) * \sum_{\bar{y}\in\mathcal{Y}} \varrho_m^s(\bar{y}) * Q_m(\bar{x}, \bar{y}) \tag{4.19}$$

Note that the above computation of the expected utility and the expected number of traders in the marketplace is also very heavy since we need to consider all possible numbers of traders choosing different actions. However, we usually calculate these when traders adopt equilibrium strategies, where only a few actions are chosen by traders (as we will show in the equilibrium analysis, see Section 4.3). Therefore, this way of calculating the marketplace's expected utility and the expected number of traders is still feasible.

4.2 The Fictitious Play Algorithm

In this section we describe how we can use fictitious play (FP) to approximate the equilibrium market selection and bidding strategies for traders in our setting. As we discussed in Section 2.1.2, the standard FP algorithm (von Neumann and Brown, 1950; Brown, 1951) is not suitable for analysing Bayesian games in which the player's type is not known to the other players. To address this, Rabinovich et al. (2009) provided a generalised fictitious play algorithm to analyse games with continuous types, a finite action space and incomplete information. However, in Rabinovich et al. (2009), researchers only showed how to use this algorithm to analyse traders' strategies in single-sided auctions. Building on this, in the following, we apply this algorithm to approximate traders' equilibrium strategies in the much more complex environment with multiple competing double auction marketplaces.

We now describe how to use the generalised FP algorithm in our setting. We first describe how to compute the best response actions against current FP beliefs. Then we 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.

We first describe how to compute the best response actions against current FP beliefs. Previously, we used Ω^b and Ω^s to denote the probability distributions of buyers' and sellers' actions respectively. In the FP algorithm, we use them to represent FP beliefs about the buyers'and sellers' actions respectively. Then, given their beliefs, we compute the buyers' and the sellers' best response functions. In the following, we describe how to compute the buyers' best response function σ^{b*}, where $\sigma^{b*}(\theta^b, \Omega^b, \Omega^s) = argmax_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$ is the best response action of the buyer with type θ^b against FP beliefs Ω^b and Ω^s. The optimal utility that a buyer with type θ^b can achieve is $\tilde{U}^*(\theta^b, \Omega^b, \Omega^s) = max_{\delta^b \in \Delta} \tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$. Considering the equations to calculate the buyer's expected utility in Section 4.1.2, we note that the buyer's expected utility $\tilde{U}(\theta^b, \delta^b, \Omega^b, \Omega^s)$ is linear in its type θ^b for a given action. 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^b_1, δ^b_2, δ^b_3 and $\ominus$, is given in Figure 4.1. Given each action, the buyer's expected utility with respect to its type is shown by $line1$, $line2$, $line3$ and $line_{\ominus}$ (i.e. x-axis) respectively. The optimal utility achieved by the buyer is represented by the set of thick piecewise linear segments. Each line segment corresponds to a type interval, where the best response action of each type in this interval is the same. In this figure, the best response action δ^b_i corresponds to the interval Ψ^b_i ($i = 1, 2, 3$) and the best response action $\ominus$ corresponds to $\Psi^b_{\ominus}$. More generally, we can create the set of distinct intervals I^b, which constitute the continuous type space of buyers, i.e. $\bigcup_{\Psi^b \in I^b} \Psi^b = [0, 1]$, which satisfy the following conditions:

- For any interval Ψ^b, if $\theta^b_1, \theta^b_2 \in \Psi^b$, then $\sigma^{b*}(\theta^b_1, \Omega^b, \Omega^s) = \sigma^{b*}(\theta^b_2, \Omega^b, \Omega^s)$, i.e. types in the same interval have the same best response action.
- For any distinct $\Psi^b_1, \Psi^b_2 \in I^b$, if $\theta^b_1 \in \Psi^b_1$, $\theta^b_2 \in \Psi^b_2$, then $\sigma^{b*}(\theta^b_1, \Omega^b, \Omega^s) \neq \sigma^{b*}(\theta^b_2, \Omega^b, \Omega^s)$

Now we have computed the best response function and also provided the set of intervals of types corresponding to the best response actions. Based on this, we can calculate the best response action distribution of buyers, which is done as follows. We know that given the buyers' type distribution function F^b and probability density function f^b, the probability that the buyer has the type in the interval Ψ^b is $\int_{\Psi^b} f(x)dx$, denoted by $F^b(\Psi^b)$. When the best response action corresponding to the interval Ψ^b_i is δ^{b*}_i, the probability that the action δ^{b*}_i is used by buyers is $\omega^b_i = F^b(\Psi^b_i)$. By calculating the probability of each action being used, we obtain the current best response action distribution of buyers, denoted by Ω^b_{br}, which is against current FP beliefs.

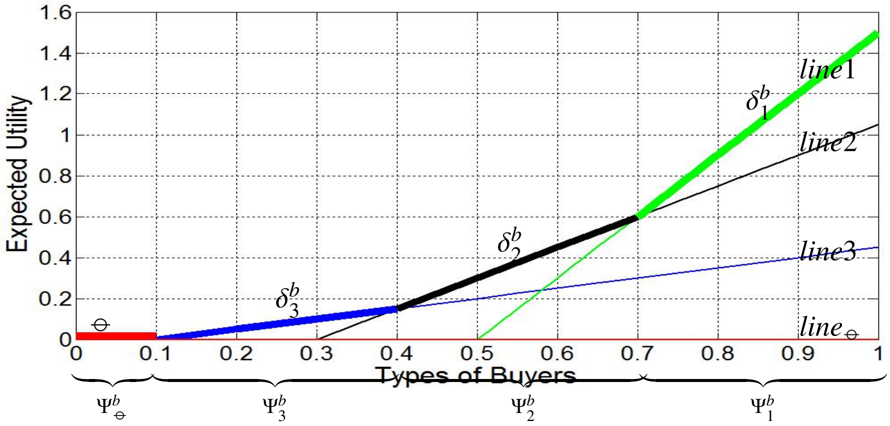

FIGURE 4.1: Piecewise linear expected utility functions.

We can then update the FP beliefs of buyers' actions, which is given by:

$$\Omega^b_{\tau+1} = \frac{\tau}{\tau+1} * \Omega^b_\tau + \frac{1}{\tau+1} * \Omega^b_{br}$$

where $\Omega^b_{\tau+1}$ is the updated FP beliefs of the buyers' actions for the next iteration round $\tau + 1$, Ω^b_τ is the FP beliefs on the current iteration round τ, and Ω^b_{br} is the probability distribution of best response actions against FP beliefs Ω^b_τ. This equation actually gives the FP beliefs on the current round as the average of FP beliefs of 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.

We now describe how to check the convergence of Nash equilibrium. Recall that in Section 2.1.2, we introduced that it is unrealistic to run the fictitious play algorithm for an infinite number of iteration rounds for the convergence, and therefore, it is often to run the algorithm for a limited number of rounds to derive the ϵ-Nash equilibrium. Recall that in the ϵ-Nash equilibrium, it is not possible for any player to gain more than ϵ in expected utility by unilaterally deviating from its strategy (see Section 2.1). Therefore, in our setting, if the difference between the expected utility of a buyer(seller) in current best response action distributions and its expected utility of adopting best response action against current best response action distributions is not greater than ϵ, the FP algorithm stops the iteration process, and the current best response actions with corresponding type intervals constitute an ϵ-Bayes-Nash equilibrium. Specifically, in our work we set ϵ as $\epsilon = 0.00001$. Formally, the measure of convergence is given by:

$$|\tilde{U}^b(\Omega^b_{br}, \Omega^s_{br}) - \tilde{U}^b_{br}(\Omega^b_{br}, \Omega^s_{br})| \leq \epsilon \; and \; |\tilde{U}^s(\Omega^b_{br}, \Omega^s_{br}) - \tilde{U}^s_{br}(\Omega^b_{br}, \Omega^s_{br})| \leq \epsilon$$

$\tilde{U}^b(\Omega^b_{br}, \Omega^s_{br})$ is the expected utility of a buyer in the best response action distributions Ω^b_{br} and Ω^s_{br}:

$$\tilde{U}^b(\Omega^b_{br}, \Omega^s_{br}) = \int_0^1 f^b(x) * \tilde{U}^b(x, \delta^b, \Omega^b_{br}, \Omega^s_{br})dx \qquad (4.20)$$

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^b_τ and Ω^s_τ, i.e. $\sigma^{b*}(x, \Omega^b_\tau, \Omega^s_\tau)$). $\tilde{U}^b_{br}(\Omega^b_{br}, \Omega^s_{br})$ is the expected utility of a buyer adopting the best response action against the current best response action distributions Ω^b_{br} and Ω^s_{br}:

$$\tilde{U}^b_{br}(\Omega^b_{br}, \Omega^s_{br}) = \int_0^1 f^b(x) * \tilde{U}^b(x, \delta^{b*}, \Omega^b_{br}, \Omega^s_{br}) dx \tag{4.21}$$

where $\delta^{b*} = \sigma^{b*}(x, \Omega^b_{br}, \Omega^s_{br})$ is the best response action of the buyer with type x against action distributions Ω^b_{br} and Ω^s_{br}. The equations for sellers are analogous.

Finally, given the calculation of best response actions, the update of FP beliefs and the measure of convergence, Figure 4.2 shows the structure of the entire FP algorithm.

Initial:
  set iteration count $\tau = 0$
  set the initial beliefs Ω^b_0 and Ω^s_0
1. **loop**
    Compute best response functions $\sigma^{b*}(\theta^b, \Omega^b_\tau, \Omega^s_\tau)$ and $\sigma^{s*}(\theta^s, \Omega^b_\tau, \Omega^s_\tau)$ against the action distribution Ω^b_τ and Ω^s_τ;
    Generate the interval Ψ^b_i corresponding to the best response action δ^{b*}_i;
    Generate the interval Ψ^s_i corresponding to the best response action δ^{s*}_i;
3. Compute current best response action distribution of buyers and sellers:
    $\Omega^b_{br} = (\omega^b_1, ..., \omega^b_{|\Delta|})$, where $\omega^b_i = F^b(\Psi^b_i)$, $i = 1, ..., |\Delta|$
    $\Omega^s_{br} = (\omega^s_1, ..., \omega^s_{|\Delta|})$, where $\omega^s_i = F^s(\Psi^s_i)$, $i = 1, ..., |\Delta|$
4. Update beliefs:
    $\Omega^b_{\tau+1} = \frac{\tau}{\tau+1} * \Omega^b_\tau + \frac{1}{\tau+1} * \Omega^b_{br}$
    $\Omega^s_{\tau+1} = \frac{\tau}{\tau+1} * \Omega^s_\tau + \frac{1}{\tau+1} * \Omega^s_{br}$
5. Measure the convergence, **if** (so), **then**
6. **return** the best response actions δ^{b*}_i and δ^{s*}_i with corresponding type intervals Ψ^b_i and Ψ^s_i
7. **end if**
8. Set $\tau = \tau + 1$
9. **end loop**

FIGURE 4.2: The fictitious play algorithm.

4.3 Equilibrium Analysis of Market Selection and Bidding Strategy

In this section, we will use the FP algorithm to analyse the traders' equilibrium strategies. We first analyse traders' equilibrium bidding strategies in a single marketplace. This will help us to understand the analysis in the more complex setting with multiple competing marketplaces. In the following, for illustrative purposes, we show our results in a specific setting with 5 buyers and 5 sellers, and 11 allowable bids(asks) unless mentioned otherwise[4]. Furthermore, we assume that the small cost for traders to enter a marketplace is set to $\varepsilon = 0.0001$. 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 marketplaces have no bias in favour of buyers or sellers when

[4]We also tried other settings. However, we still obtained the similar results.

allocating surpluses. Finally, we assume that both buyers and sellers' types are independently drawn from a uniform distribution.

4.3.1 A Single Marketplace

As we introduced in Section 2.3.2.1, many heuristic bidding strategies for double auctions have been proposed in the literature (GD, ZIP, ZI-C and so on). However, they all fail to answer what exactly traders should bid in equilibrium. This is important since how traders bid in a given marketplace will affect their expected utilities and this, in turn, their selection of marketplaces. Furthermore, market fees may also affect traders' bidding strategies. Therefore, in this section, we first analyse the equilibrium bidding strategies of traders in a single marketplace without fees, and then analyse the strategies when the marketplace charges different types of fees.

We first consider the case where the marketplace charges no fees to traders. We use the FP algorithm to analyse the traders' equilibrium strategies, and find that starting from different initial beliefs of traders' actions, all traders who choose the marketplace eventually converge to the same pure Nash equilibrium bidding strategy, which is shown in Figure 4.3. The gray line represents buyers' bids in equilibrium and the black line represents sellers' asks in equilibrium. From this figure, we can see what traders will bid corresponding to their types in equilibrium. We find that buyers shade their bids by decreasing their bids, and sellers shade their asks by increasing their asks, 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 ε.

Now we consider how different types of fees (registration, transaction, profit and transaction price percentage fees) can affect the traders' equilibrium bidding strategies. First we consider that the marketplace charges a registration fee. For example, when we assume that it charges 0.1 registration fee, then the traders' equilibrium bidding strategies are shown in Figure 4.4. We can see that, compared to the case where no fees are charged (see Figure 4.3), there exists a bigger range of types of traders not choosing this marketplace. This is because the registration fee causes negative profits for them. In addition, we further find that when registration fees are charged, both buyers and sellers will shade less (compared to Figure 4.3) in order to increase the probability of being matched.

After analysing the traders' equilibrium bidding strategies in the marketplace charging a registration fee, we now consider the case that the marketplace charges a transaction fee. As an example, we assume that the marketplace charges 0.1 transaction fee. The results are shown in Figure 4.5. We find that compared to the case that no fees are charged (see Figure 4.3), there exists a bigger range of types of traders not choosing the marketplace. However, compared to the case that a registration fee is charged (see Figure 4.4), more traders are willing to submit shouts since the payment only happens after they make transactions.

Then we consider the case that the marketplace charges a profit fee. As an example, we assume

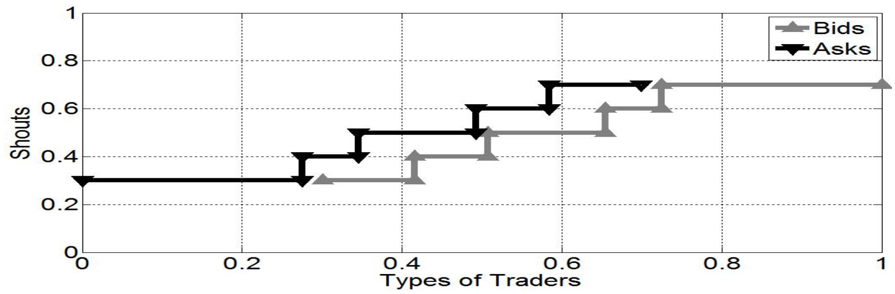

FIGURE 4.3: Equilibrium strategies with the same number of buyers and sellers.

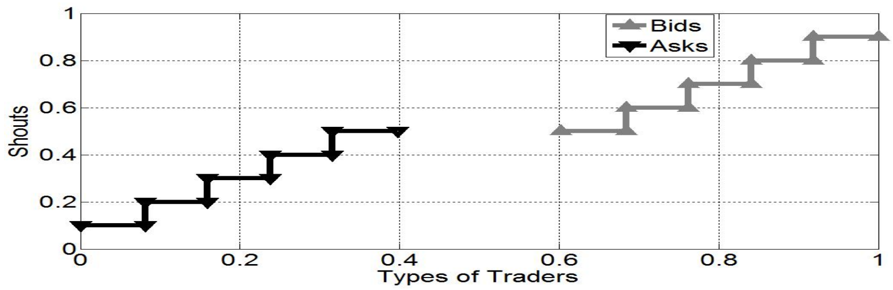

FIGURE 4.4: Equilibrium strategies of traders with registration fees.

that the marketplace charges 50% profit fee. The results are shown in Figure 4.6. We still find that compared to the case that no fees are charged (see Figure 4.3), there exists a bigger range of types of traders not choosing the marketplace. However, compared to the case that a transaction fee is charged (see Figure 4.4), more traders are willing to submit shouts. This is because the profit fee is a percentage fee charged on the observed trading surplus (which is the difference of the matched bid and ask), and when traders shade their shouts, a profit fee will not cause negative profits for them. Furthermore, we find that charging profit fees causes traders to shade their shouts more. This is because when a profit fee is charged, the marketplace extracts profits from traders according to their trading surpluses, which are the differences of matched bids and asks. Then in order to keep more profits, the traders try to reduce the trading surpluses by lowering bids or increasing asks.

Now we consider the case that the marketplace charges a transaction price percentage fee. As an example, we assume that the marketplace charges 20% transaction price percentage fee. The results are shown in Figure 4.7. We find that compared to the case where no fees are charged (see Figure 4.3), there exists a bigger range of types of traders not choosing the marketplace. Furthermore, we find that in this case, in contrast to the above cases, the behaviour of buyers and sellers are not symmetric. Sellers shade their asks less than buyers. For example, when types of sellers are within [0.1, 0.19], sellers submit asks 0.3 (i.e. the shading is from 0.11 to 0.2), compared to that when types of buyers are within [0.81 0.9], buyers submit bids 0.6 (i.e. the shading is from 0.21 to 0.3). The reason is as follows. The transaction price is at the middle of the shouts of the matched buyer and seller. When a transaction price percentage fee is charged, in order to reduce the payment, buyers will shade to decrease their bids in order to reduce the transaction prices. However, for sellers, on one hand, they want to shade more by increasing

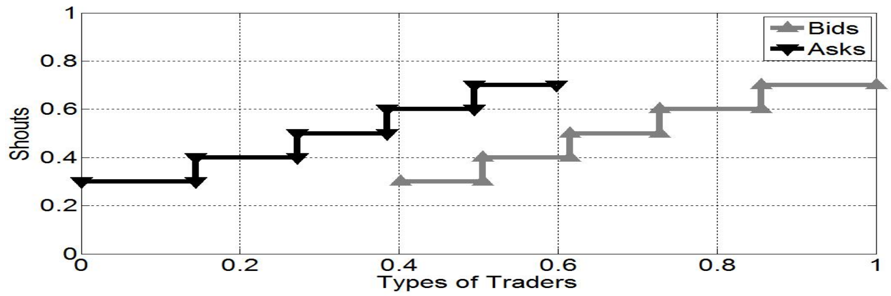

FIGURE 4.5: Equilibrium strategies of traders with transaction fees.

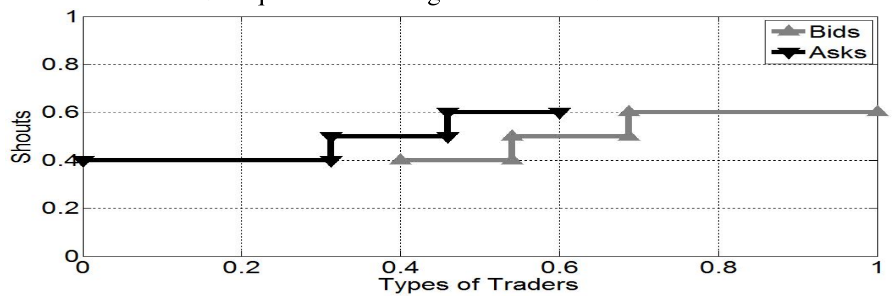

FIGURE 4.6: Equilibrium strategies of traders with profit fees.

their asks; on the other hand, they want to shade less to decrease transaction prices to reduce the payment. They have to make a trade-off. That's why in this case, buyers shade their bids more in contrast to that sellers shade their asks less.

Note that in the above analysis, we choose a specific value for each type of fees. We also try other values of fees. We find that higher fees will cause more traders not to participate in the marketplace. However, for the bidding behaviour of traders still participating in the marketplace, the conclusions obtained in the above analysis are still applied. For example, when a higher registration fee is charged, traders have to shade their shouts less in order to increase the probability of being matched; when a higher transaction price percentage fee is charged, sellers will shade much less than buyers; and when a higher profit fee is charged, traders will shade their shouts more in order to keep more profits.

In the above, we considered the case with the same number of buyers and sellers. In often happens that there are different numbers of buyers and sellers in the marketplace. In this situation, buyers and sellers have different market power, which will affect their bidding strategies. We now analyse this issue in our setting. For example, we consider the case that there are 8 buyers and 5 sellers and no fees are charged to traders. Then the traders' equilibrium bidding strategies are shown in Figure 4.8. Compared to Figure 4.3, as can be expected, we see that because there are more buyers than sellers, the competition between buyers is more severe, and thus they have to raise their bids. For sellers, since they have a higher probability of being matched, they raise their asks. At this moment, sellers have more market power, and can therefore extract more profit from their transactions.

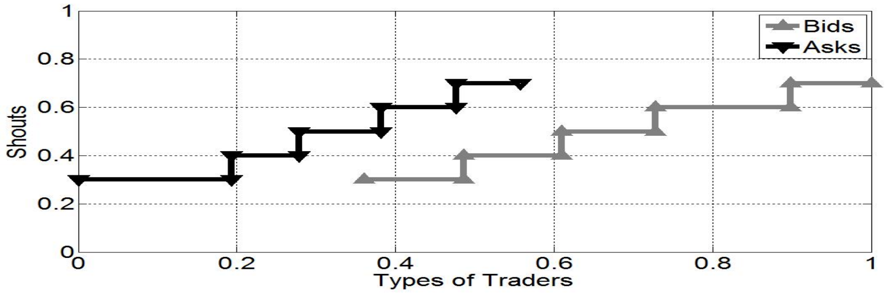

FIGURE 4.7: Equilibrium strategies of traders with transaction price percentage fees.

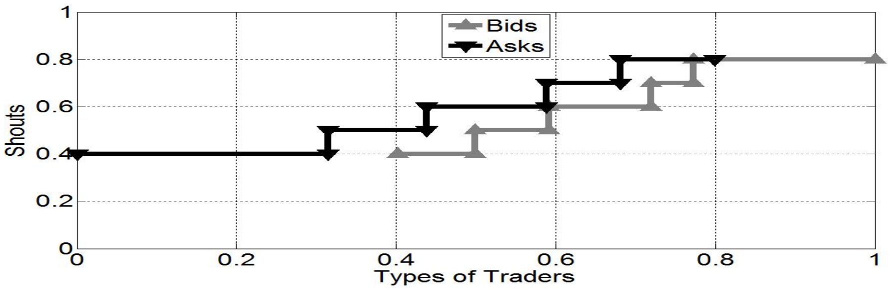

FIGURE 4.8: Equilibrium strategies with 8 buyers and 5 sellers.

4.3.2 Competing Marketplaces

In the above, we analysed how traders submit shouts in a single marketplace in equilibrium. This analysis will provide foundations for analysing how traders submit shouts in the environment with competing marketplaces. As we will show, in the competing marketplace context, when traders converge to one marketplace, their bidding strategies are exactly the same as that in a single marketplace environment. Based on this, we now analyse traders' equilibrium market selection and bidding strategies in the competing marketplace environment with two competing marketplaces[5]. As mentioned in Section 1.2, there are three types of trading environments: single-home, multi-home and hybrid trading environments. Furthermore, when multiple goods are traded by the same trader across multiple marketplaces, 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 good properties. However, in the multi-home and hybrid trading environments where multiple goods can be traded by each trader, we need to consider these different good properties. In the following, we first analyse traders' equilibrium strategies in the single-home trading environment. Then we extend the analysis to the multi-home and hybrid trading environments with different good properties.

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

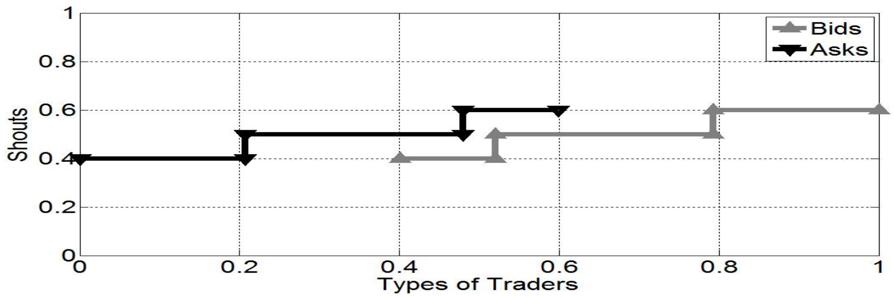

FIGURE 4.9: Equilibrium strategies of traders with different types of fees.

4.3.2.1 Single-Home Trading

We now consider traders' equilibrium strategies in the single-home trading environment, where traders can only enter one marketplace at a time. In this environment, traders can only trade one unit of good, and thus its value on the good is its type, i.e. $\alpha_1^b = \alpha_1^s = 1$. In order to understand the effects of fees on the traders' strategies, we first consider the basic setting where the marketplaces charge no fees to traders, before we consider fees. 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 at this moment, the traders will eventually converge to marketplace 1 or 2 with the same probability. In addition, we find that the traders' bidding strategies in the converged marketplace are the same as the case with a single marketplace (i.e. Figure 4.3).

We now consider what happens to the traders' behaviour when they are charged fees. First we consider the cases where both marketplaces charge the same type of fees (registration, transaction, profit or transaction price percentage). We run simulations with many possible initial beliefs, and find that traders eventually converge to one marketplace, which depends on initial FP beliefs and market fees. Since traders converge to one marketplace, the equilibrium bidding strategies are the same as the case with a single marketplace (i.e. Section 4.3.1).

Now we consider the cases where marketplaces charge different types of fees. Firstly, we consider that marketplace 1 charges a profit fee and marketplace 2 charges a registration fee. For example, we consider that marketplace 1 charges a relatively high profit fee of 90%, and marketplace 2 charges a registration fee of 0.1. If initial beliefs are uniform (i.e. all actions are equally probable), we find that all traders eventually converge to marketplace 1, and the equilibrium bidding strategies are shown in Figure 4.9. The reason for the traders converging to marketplace 1, which seems more expensive, is as follows. When a high profit fee is charged, the traders shade their shouts more to keep profits (as can be seen in Figure 4.9). However, shading has no effect in the case of registration fees. Therefore, traders prefer the marketplace charging profit fees compared to registration fees. Furthermore, we ran simulations with many other fee combinations and different initial beliefs, and always find that all traders converge to one marketplace. This is different from the analysis in Section 3.2.3.2, where when different types of fees are charged, traders may converge to different marketplaces in equilibrium (rich traders converge to the marketplace charging a registration fee, and poor traders converge to the marketplace charging a profit fee). The reason is that traders can hide their actual trading surpluses by shading their shouts, and thus even when a high profit fee is charged, traders still can keep profits. As per the analysis in Chapter 3, the positive size effect has a larger impact than the negative size effect, and all traders will converge to one marketplace.

Furthermore, we consider the cases of other combinations of different types of fees, such as marketplace 1 charging a profit fee, and marketplace 2 charging a transaction price percentage fee, or marketplace 1 charging a registration fee, and marketplace 2 charging a transaction price percentage fee, and so on. However, we still find that traders eventually converge to one marketplace, and the equilibrium bidding strategies for traders is the same as the case of a single marketplace.

4.3.2.2 Multi-Home Trading

So far we have analysed equilibrium strategies for traders in the single-home trading environment. Now we extend the analysis to the multi-home trading environment where both buyers and sellers can enter multiple marketplaces at a time. In such an environment, multiple goods may be traded by each trader. Therefore we now need to consider the properties of goods since these will affect traders' strategies. Specifically, here we assume that for sellers, all goods are independent (i.e. when obtaining one unit of good, $\alpha_1^s = 1$, and when obtaining two units of goods, $\alpha_2^s = 2$) since a seller's value on sold goods is usually equal to the sum of its value on each individual sold good (i.e. additive). Thus in the multi-home trading 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 traders' equilibrium strategies by considering different good properties for buyers. Furthermore, since we focus on the effects of different trading environments and good properties on the traders' equilibrium strategies, we assume that both competing marketplaces charge no fees (the analysis of the cases with fees is similar).

Independent Goods:

In the multi-home trading environment with independent goods, buyers' values on goods are additive, i.e. the good property coefficient $\alpha_1^b = 1$, $\alpha_2^b = 2$. In such an environment, buyers and sellers will enter both marketplaces when their expected utilities (equal to expected values minus the expected payments) in both marketplaces are positive. The analysis of their bidding strategies is then identical to the single marketplace setting (i.e. Section 4.3.1).

Substitutable Goods:

Now we analyse the equilibrium strategies for traders in the multi-home trading environment with substitutable goods for buyers (i.e. $\alpha_1^b = 1$, and $1 \leq \alpha_2^b < 2$). Here we consider perfectly substitutable goods, i.e. $\alpha_1^b = 1$ and $\alpha_2^b = 1$. This means that when the buyer with type θ^b

wins one good, it obtains value θ^b and pays for the good; and when this buyer wins two goods, it only obtains value θ^b, but pays for two goods. When both 5 buyers and 5 sellers are multi-home trading, totally, there are $2 * 12^2 = 288$ different actions that traders can take, in contrast to $2 * 12 * 2 = 48$ different actions in the single-home environment. Since sellers are multi-home trading with independent goods and both competing marketplaces are identical, in order to reduce the action space, we simplify the analysis by assuming that 5 sellers only submit asks in marketplace 1, and another 5 sellers only submit asks in marketplace 2. Now the total number of actions for traders is reduced from 288 to $12^2 + 24 = 168$. The results is shown in Figure 4.10(a) and 4.10(b). We can find that, in this setting, the same buyer (or buyers with the same type) will bid differently in two marketplaces. The reason is as follows. Because of perfectly substitutable goods, buyers would like to purchase and pay for only one good. Therefore, buyers with high types will choose to only bid in one marketplace, and will 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, for buyers with lower types, in order to make transactions, they have to bid in both marketplaces in order to increase the probability of being matched. In this situation, they may make transactions in both marketplaces, which means that they have to pay for two goods. Then in order to make transactions and decrease the payment, they may choose to bid slightly higher in one marketplace, and lower in another one. That's why buyers bid differently in two marketplaces. Furthermore, we note that some of these buyers with lower types will bid higher in one marketplace, and the other will bid higher in another one. This is because all buyers bidding higher in one marketplace results in fierce competition between them, and thus some of them will choose to bid higher in another one in order to increase the probability of being matched.

Complementary Goods:

Now we analyse the setting with complementary goods for buyers. We first consider the case with perfectly complementary goods, i.e. $\alpha_1^b = 0$ and $\alpha_2^b = 1$. This means that when the buyer with type θ^b wins one good, it obtains zero value, but needs to pay for the good; and when this buyer wins two goods, it obtains value θ^b, and pays for two goods. Furthermore, in order to reduce the action space, we still assume that 5 sellers only submit asks in marketplace 1, and another 5 sellers only submit asks in marketplace 2. The results are shown in Figures 4.11(a) and 4.11(b). We can see that with perfectly complementary goods, only buyers with high types will bid in marketplaces, and they will bid in both marketplaces (if they only bid in one marketplace and make a transaction, they need to pay for the good, but obtain zero value because of the perfectly complementary goods). Buyers with relatively low types will not enter marketplaces since even though they bid in both marketplaces, because of their low types and perfectly complementary goods, they are more likely to obtain negative profits. Furthermore, we find that buyers shade their bids more. For example, when buyers' types are within [0.825, 1.0], buyers will bid 0.5 in both marketplaces. The reason is as follows. With perfectly complementary goods, the buyer with type θ^b purchasing two goods, only obtains value θ^b, but has to

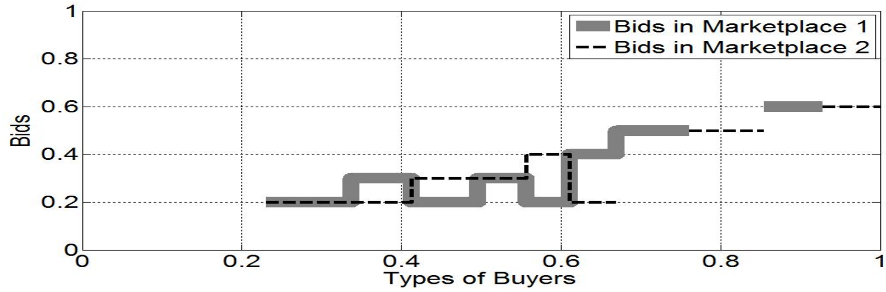

(a) Buyers

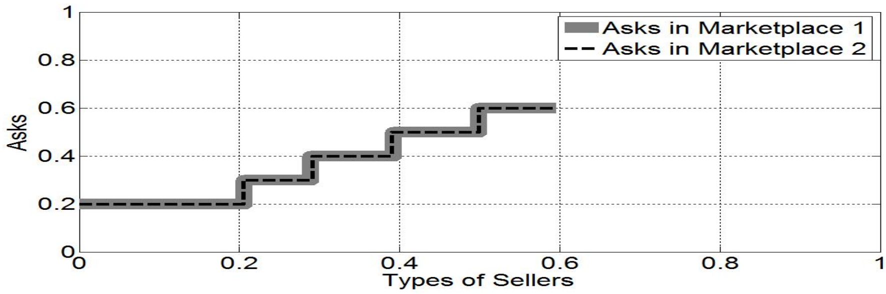

(b) Sellers

FIGURE 4.10: Equilibrium strategies of traders in the multi-home trading environment with perfectly substitutable goods for buyers.

pay for two goods. In order to keep profits, buyers should shade their bids more. Furthermore, we extend this analysis to another case that $\alpha_1^b = 1$ and $\alpha_2^b = 4$, i.e. when the buyer with type θ^b wins one good, it obtains value θ^b and pays for the good; and when this buyer wins two goods, it obtains value $4 * \theta^b$ (four times of its type), but only pays for two goods. The results are shown in Figures 4.12(a) and 4.12(b). Compared to the perfectly complementary goods, we find that buyers still bid in two marketplaces. However, they will 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, buyers' values for goods are very high (i.e. $4 * \theta^b$). In this situation, in terms of obtaining more utility, it is better for the buyers to increase the bids to increase the probability of making transactions than to shade bids. Furthermore, this will cause a bigger range of sellers to ask since when buyers increase their bids, sellers with high types still can make transactions.

4.3.2.3 Hybrid Trading

Now we analyse the equilibrium strategies for traders in the hybrid trading environment where one side can only participate in one marketplace, and the other side can participate in multiple marketplaces. Specifically, we consider the case that buyers can participate in multiple marketplaces and sellers can only participate in one marketplace at the same time. Note that the results of the opposite case where sellers can choose multiple marketplaces and buyers can only choose one are similar.

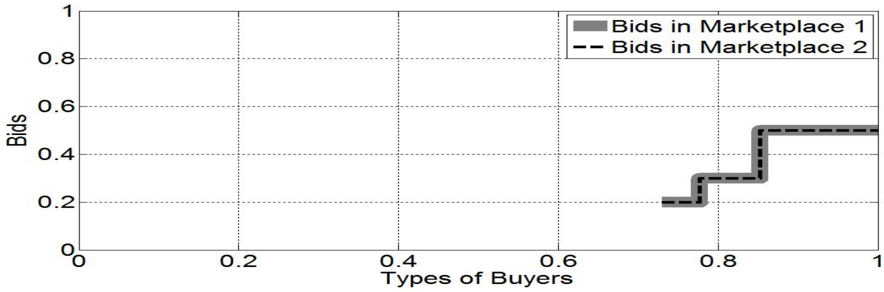

(a) Buyers

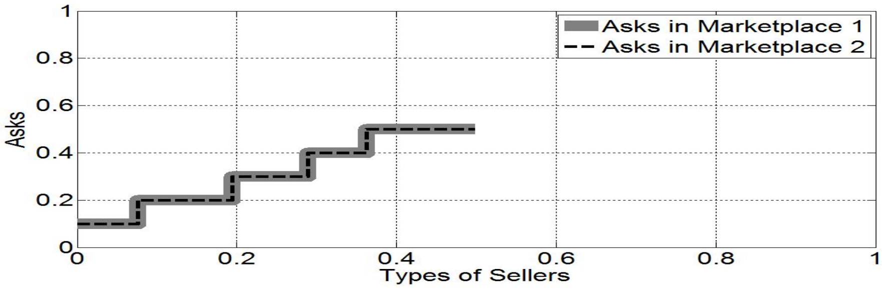

(b) Sellers

FIGURE 4.11: Equilibrium strategies of traders in the multi-home trading environment with perfectly complementary goods for buyers.

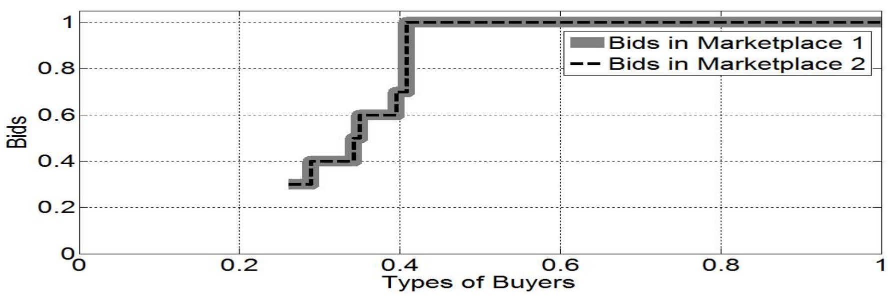

(a) Buyers

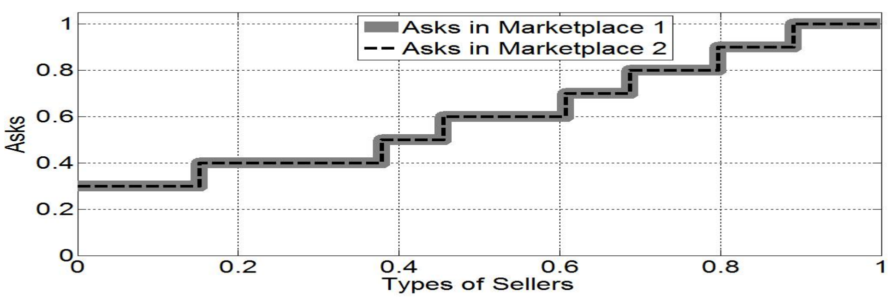

(b) Sellers

FIGURE 4.12: Equilibrium strategies of traders in the multi-home trading environment with complementary goods for buyers.

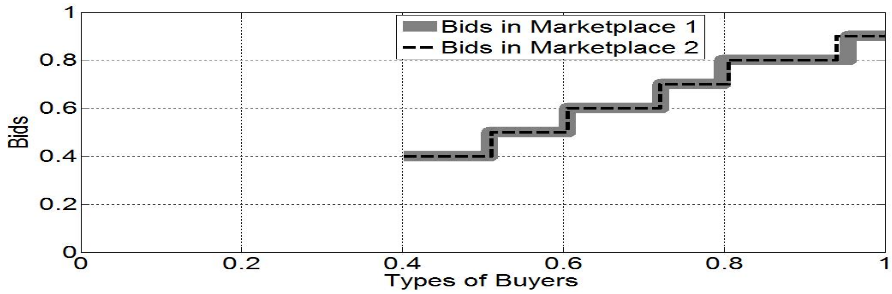

(a) Buyers

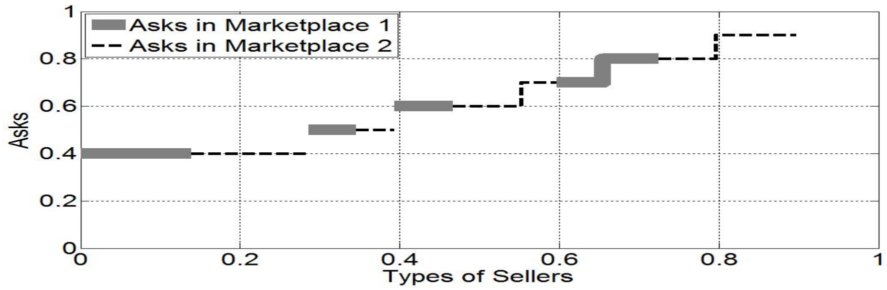

(b) Sellers

FIGURE 4.13: Equilibrium strategies of traders in the hybrid trading environment with independent goods for buyers.

Independent Goods:

We now analyse the case with independent goods for buyers (i.e. $\alpha_1^b = 1$, $\alpha_2^b = 2$). The results are shown in Figures 4.13(a) and 4.13(b). From these figures, we can see that sellers eventually split and place asks in different marketplaces in equilibrium. In this situation, two competing marketplaces co-exist. This co-existence is caused by the negative size effect. In more detail, the sellers have to compete with each other in order to be matched with buyers and make transactions, and thus they prefer those marketplaces with fewer sellers. Because identical buyers stay in both competing marketplaces in this case, then the attractiveness from the buyers to the sellers (i.e. the positive size effect) in both marketplaces is the same. At this moment, the internal competition between the sellers (i.e. the negative size effect) takes effect, which drives the sellers to stay in different marketplaces. Another interesting phenomenon is that compared to the traders' equilibrium strategies in Figure 4.3, we find that buyers raise their bids and sellers also raise their asks in this case. The reason is as follows. As the sellers are split in two marketplaces, then in each marketplace the number of sellers is less than the number of buyers. Thus as per our previous analysis (see Figure 4.8), sellers have more market power than buyers, and so buyers raise their bids in order to be matched and sellers raise their asks to extract more profits from transactions.

Substitutable Goods:

Now we analyse the cases when goods are substitutable for buyers. We first analyse the case with perfectly substitutable goods (i.e. $\alpha_1^b = 1$, $\alpha_2^b = 1$). In this case, we find that in equilibrium traders will only choose one marketplace, and the bidding strategies are the same as the case with

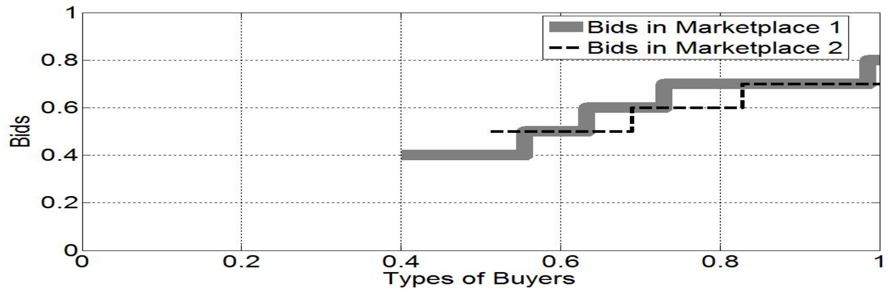

(a) Buyers

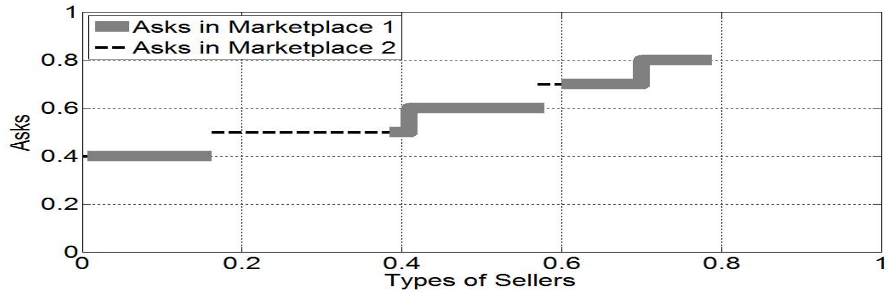

(b) Sellers

FIGURE 4.14: Equilibrium strategies of traders in the hybrid trading environment with substitutable goods for buyers.

a single marketplace charging no fees (see Figure 4.3). The reason is as follows. When goods are perfectly substitutable for buyers, buying more than one good does not mean obtaining more value for buyers, but does mean paying more. Thus buyers will prefer to bid in one marketplace. This will cause sellers to converge to that marketplace. Furthermore, since in the above analysis when buyers' expected values on goods are additive, we find that buyers will bid in multiple marketplaces. Thus we hypothesise that as buyers' values on multiple goods increase, they will begin to prefer to bid in multiple marketplaces, and sellers may split in two marketplaces because of the negative size effect. This is confirmed by running experiments with different values of α_2^b. We find that when $\alpha_2^b \geq 1.8$, buyers will begin to bid in two marketplaces. Specifically, when $\alpha_2^b = 1.8$, the traders' equilibrium bidding strategies are shown in Figures 4.14(a) and 4.14(b), from which we can see that buyers bid in two marketplaces and sellers distribute in two marketplaces.

Complementary Goods:

Finally, we analyse the case with complementary goods for buyers in the hybrid trading environment. Firstly, we consider perfectly complementary goods for buyers (i.e. $\alpha_1^b = 0$ and $\alpha_2^b = 1$). When there are 5 buyers and 5 sellers, we find that in equilibrium, no traders will 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 more in order to make positive profits. However, since sellers are single-home trading, sellers have more market power when they split in two marketplaces, and buyers have to increase their bids in order to increase the probability of being matched. Now buyers cannot shade at a high degree. When they purchase two goods,

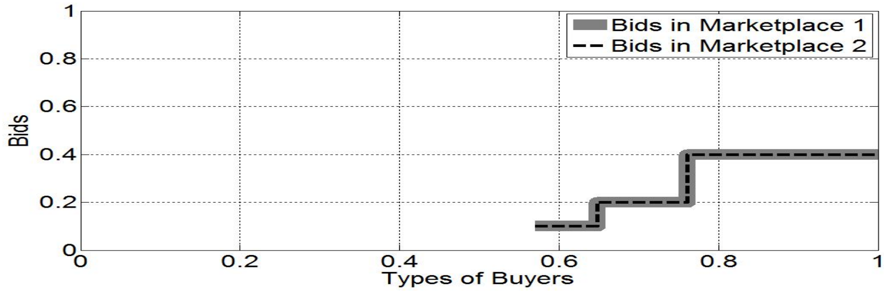

(a) Buyers

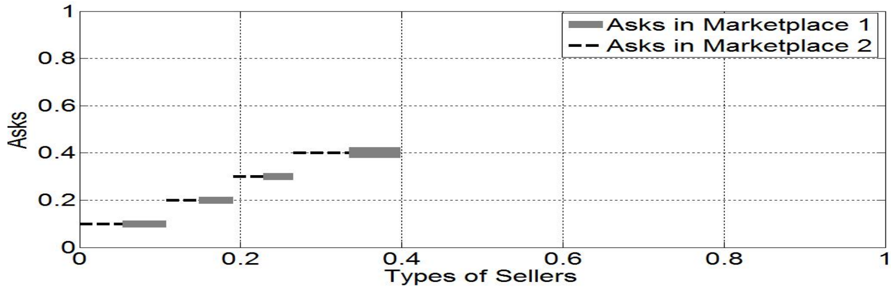

(b) Sellers

FIGURE 4.15: Equilibrium strategies of traders in the hybrid trading environment with perfectly complementary goods for buyers.

their values of the two goods are equal to their types, but have to pay for two goods. In this situation, buyers may have negative profits. Now buyers will not choose any marketplaces, and then sellers will also not choose any marketplaces. If we increase the number of sellers (to decrease their market power) and decrease the number of buyers (to increase their market power), for example, we consider 2 buyers and 10 sellers, the results are shown in Figures 4.15(a) and 4.15(b). We can see that buyers will bid in two marketplaces, and shade more to keep profits. Because of the negative size effects, sellers will distribute in two marketplaces. Furthermore, in the case with 5 buyers and 5 sellers, we change the property coefficient as $\alpha_1^b = 1$ and $\alpha_2^b = 4$ (i.e. when purchasing two goods, the buyer with type θ^b can obtain value $4 * \theta^b$). The results are shown in Figures 4.16(a) and 4.16(b). We can see that buyers will bid in two marketplaces, and instead of shading, they increase their bids to increase the probabilities of making transactions, since when making transactions in both marketplaces, they will obtain very high values (i.e. four times their types). For sellers, since buyers' bids in both marketplaces are identical, each type of seller has equal probability to ask in each marketplace. For example, when sellers' types are within [0, 0.257], sellers can ask 0.6 in marketplace 1 or ask 0.6 in marketplace 2. Furthermore, since buyers increase their bids, sellers increase their asks to make more profits.

4.4 Equilibrium Analysis of Charging Strategies

In Section 4.3.1, we analysed the traders' equilibrium bidding strategies in a single marketplace with different types of fees charged. From the analysis, we find that when registration fees or

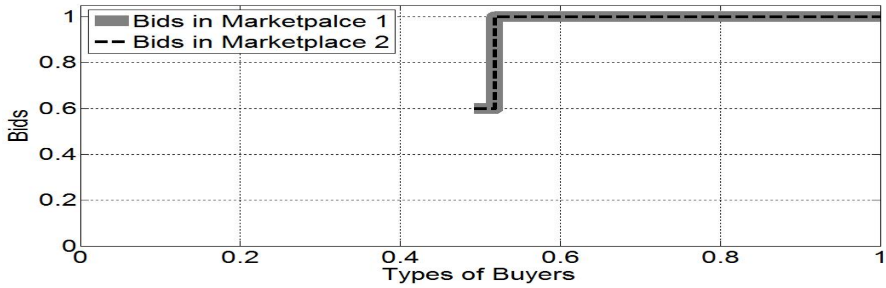

(a) Buyers

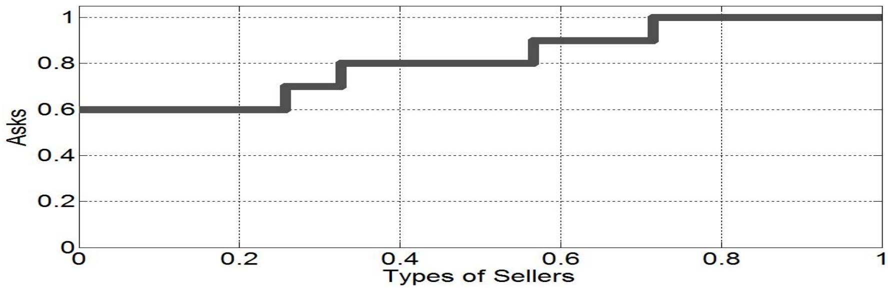

(b) Sellers

FIGURE 4.16: Equilibrium strategies of traders in the hybrid trading environment with complementary goods for buyers.

transaction fees are charged, traders may choose to leave the marketplace; when profit fees are charged, traders will shade their shouts more, and high profit fees cannot guarantee high market profits; and when transaction price percentage fees are charged, compared to buyers, sellers shade less in order to reduce transaction prices, and thus reduce the payments. From all of this, we can see that it is not obvious which type of fees is the most effective in terms of maximising market profits and keeping traders. Given this, in the following, we first analyse this problem in a single marketplace, which will help us find the most effective type of fees that we can use in the competition. Then we will analyse how competing marketplaces set fees in equilibrium.

4.4.1 A Single Marketplace

As before, we analyse a single marketplace with 5 buyers and 5 sellers, but now consider the profit made by the marketplace when different types of fees are charged. We start by analysing registration fees. Specifically, we discretize registration fees from 0 to 1 with step size 0.01. For each registration fee, we use FP to analyse the equilibrium strategies for traders. Then according to the traders' action distributions in equilibrium, we use Equations 4.17 and 4.19 to calculate the expected profit for the marketplace and the expected number of traders entering the marketplace respectively. By so doing, we know for each level of registration fee, how much profit the marketplace can extract from traders and how many traders will enter the marketplace. These are shown in Figure 4.4.1, where Figure 4.17(a) shows the expected market profit and Figure 4.17(b) shows the expected number of traders for different registration fees. From these results, we can see that, as the marketplace increases its registration fee, the market profit first

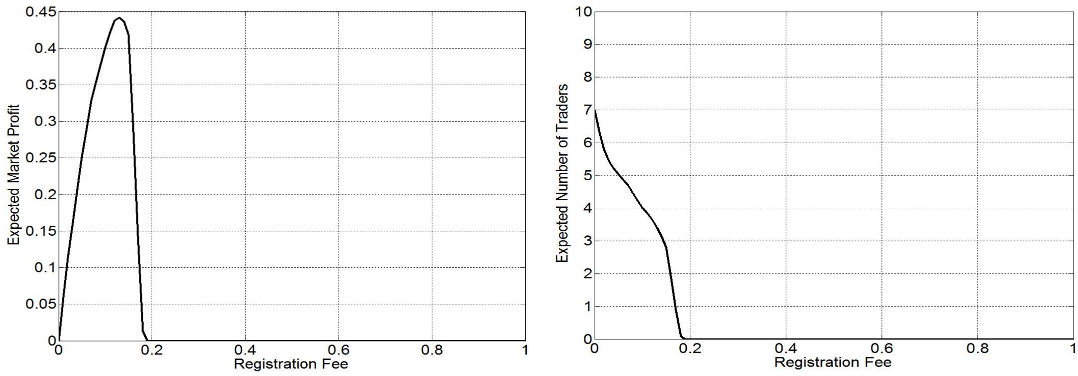

(a) Market profit (b) Number of traders

FIGURE 4.17: Market profit and number of traders in a single marketplace charging a registration fee.

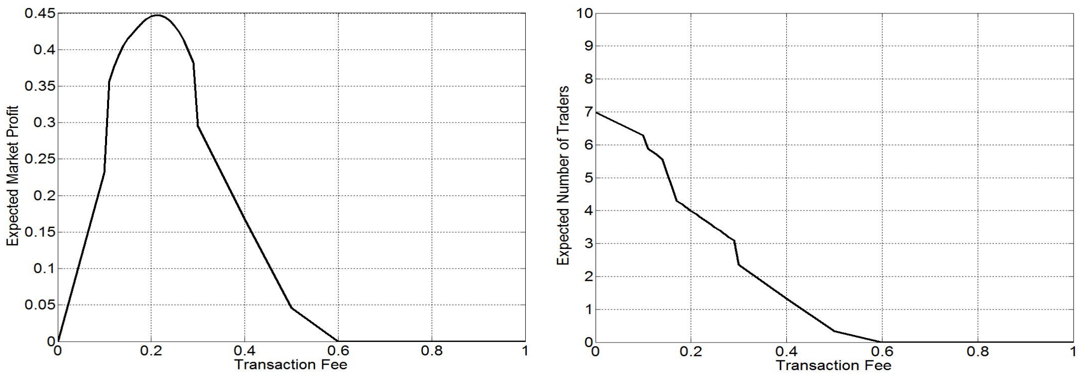

(a) Market profit (b) Number of traders

FIGURE 4.18: Market profit and number of traders in a single marketplace charging a transaction fee.

increases, and when the marketplace charges a 0.14 registration fee, it obtains the maximum market profit. From Figure 4.17(b), we can see how this relates to the number of traders. Initially when the marketplace charges no fees to traders, it contains around 7 traders (the small cost ε causes the other 3 traders not to choose the marketplace). When the marketplace increases the registration fee, we can see that traders leave the marketplace quickly. When the registration fee increases to around 0.19, no traders will stay in this marketplace. As traders leave the marketplace, from Figure 4.17(a), we can see that the market profit also decreases, and going to 0 when the registration fee increases to around 0.19.

Now we analyse how the marketplace sets its transaction fee to make a profit. The results are shown in Figures 4.18(a) and 4.18(b). From this, we can see that the marketplace obtains the maximum profit when it charges a transaction fee of around 0.22. When the transaction fee reaches 0.6, no traders stay in the marketplace.

Furthermore, we analyse how the marketplace sets its profit fee to make a profit. The results are shown in Figures 4.19(a) and 4.19(b). We can see that as the profit fee increases, the market profit first increases to the maximum point, and then decreases. Furthermore, from Figure

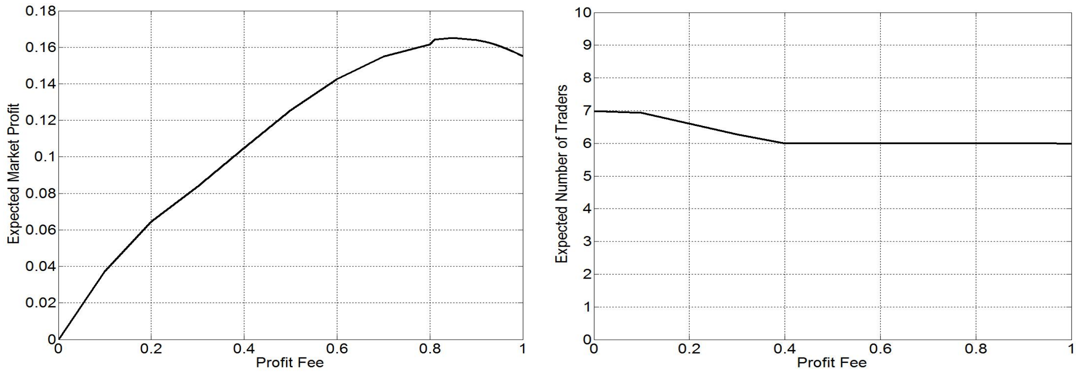

(a) Market profit (b) Number of traders

FIGURE 4.19: Market profit and number of traders in a single marketplace charging a profit fee.

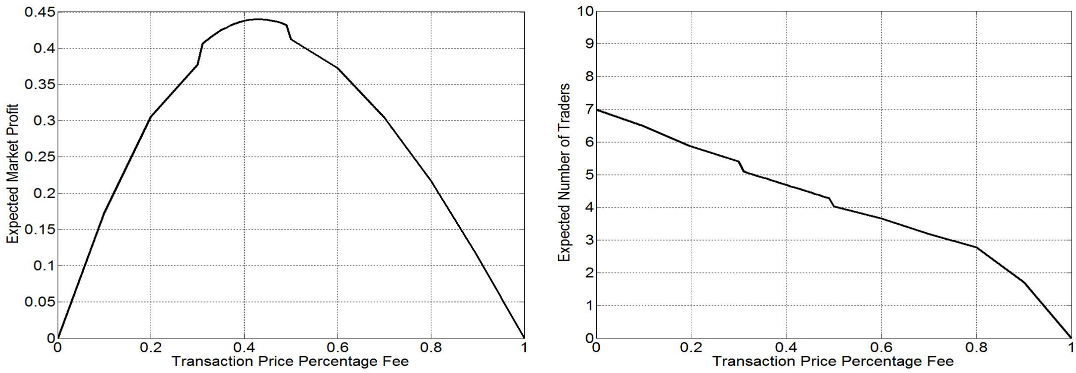

(a) Market profit (b) Number of traders

FIGURE 4.20: Market profit and number of traders in a single marketplace charging a transaction price percentage fee.

4.19(b), we can see that the marketplace can keep traders at a good level (around 6 traders staying in the marketplace even when it charges 100% profit fee), since traders can shade their shouts to hide actual trading surpluses, and thus reduce payments.

Finally, we analyse what happens when the marketplace charges a transaction price percentage fee. From Figure 4.20(a) we find that when the marketplace charges a 43% transaction price percentage fee, the maximum market profit is reached. Furthermore, from Figure 4.20(b), we can see that traders still choose the marketplace even though a high transaction price percentage fee is charged.

Now we compare the effects of different types of fees on obtaining market profits and keeping traders. From Figures 4.19(a) and 4.19(b), we can see that when a profit fee is charged to traders, the marketplace is more likely to keep traders. However, even though a very high profit fee is charged, the extracted profit is still low. From Figures 4.17(a), 4.18(a) and 4.20(a), we can see that the effects of the registration, transaction and transaction price percentage fees on obtaining market profit are similar, In more detail, when this marketplace charges a 0.14 registration fee, or a 0.22 transaction fee, or a 43% transaction price percentage fee, it will obtain the maximum profit which is around 0.44. However, at this point, only around 3.1 traders stay in the marketplace when a registration fee is charged, around 3.6 traders stay when a transaction fee is charged, and around 4.6 traders stay when a transaction price percentage fee is charged. Given this, we can conclude that the marketplace charging a transaction price percentage fee is the best on obtaining market profits and keeping traders. Furthermore, we can see that a transaction fee is better than a registration fee. Thus in the design of a charging strategy for the CAT competition (in Section 5.4), we will let the marketplace charge a transaction fee, instead of charging a registration fee. Note that although the transaction price percentage fee is the most effective, it is not allowed in the CAT competition.

4.4.2 Competing Marketplaces

In the above, we analysed how a single marketplace sets its fee to maximise profit while keeping traders. However, this analysis did not consider the competition between multiple marketplaces. Given this, in this section, we will analyse how marketplaces set fees in equilibrium when competing with each other. We analyse this in the single-home trading setting since such a trading mechanism results in a highly competitive environment (where marketplaces have to compete fiercely with each other to attract traders), and we are interested in analysing how both marketplaces set fees in such an environment[6]. In the following analysis, we discretize fees from 0 to 1 with step size 0.1. Then we obtain different fee systems. For each fee system, we calculate the marketplaces' expected utilities. In more detail, for a given fee system, we repeat the experiments by trying different initial fictitious play beliefs[7]. For each set of initial FP beliefs, we run the fictitious play algorithm and obtain the traders' equilibrium strategies. Given the equilibrium strategies of the traders, by using Equation 4.17, we then calculate the marketplaces' expected utilities for the given fee system when starting from the particular FP beliefs. When repeating the experiments from different initial FP beliefs, we obtain the average utilities of marketplaces for this given fee system. We repeat this process for different fee systems, and obtain a payoff table, from which we can analyse the equilibrium fee system for marketplaces. In the following, we analyse how marketplaces set fees in three different cases: both marketplaces only charge registration fees; only charge profit fees; and one marketplace charges a registration fee, and the other charges a profit fee (which is the same as what we did in Chapter 3). The analysis of marketplaces charging other types of fees is similar.

In the first case where both marketplaces only charge registration fees, there are 121 different fee systems. The payoff table are shown in Table 4.1, from which, by using Gambit, we can see that both marketplaces charging 0.1 registration fee constitutes a NEQ fee system.

In the second case, we assume that both marketplaces only charge profit fees. The payoff table is shown in Table 4.2. Compared to the analysis using EGT in Section 3.3.1, we find that

[6]This analysis can also be easily extended to multi-home trading and hybrid trading environments.

[7]Similar to the analysis in Section 3.2.3, we also find the existence of a lock-in region, i.e. when traders start from some FP beliefs, where they initially prefer some marketplace, then even though this marketplace charges a higher fee, traders may eventually converge to this marketplace in equilibrium.

in equilibrium, both marketplaces charge higher profit fees. This is because traders can now shade their shouts, whereas before we assumed that traders had a truth-telling bidding strategy. Specifically, both marketplaces charging a 60% profit fee with probability 0.067, a 70% profit fee with probability 0.593 and a 80% profit fee with probability 0.340 constitutes a mixed Nash equilibrium.

Now we consider the case that different types of fees are charged to traders. We consider that marketplace 1 charges a registration fee and marketplace 2 charges a profit fee. The payoff table is shown in Table 4.3. From this, we find that marketplace 1 charging a 0.1 registration fee, and marketplace 2 charging a 90% profit fee constitutes a NEQ fee system.

4.5 Summary

In this chapter, we used a FP algorithm to analyse how traders select marketplaces and submit shouts, and how competing marketplaces set fees in equilibrium in a setting with continuous trader types. Specifically, we first analysed traders' equilibrium bidding strategies in a single marketplace, where we found that traders shade their shouts in equilibrium. We further analysed the effect of different types of market fees on the traders' equilibrium bidding strategies and observed that registration fees cause a bigger range of traders to not choose the marketplace; profit fees cause traders to shade their shouts more; and transaction price percentage fees cause sellers to shade relatively less than buyers' shading. Then we analysed the traders' equilibrium market selection strategies and bidding strategies in the single-home trading environment with multiple marketplaces. We found that all traders eventually converge to bid in one marketplace. Competing marketplaces cannot co-exist even though they charge different types of fees. This is contrary to the conclusion in Chapter 3 that competing marketplaces may co-exist when one marketplace charges a registration fee and another charges a profit fee. This difference occurs because in this setting, traders can reduce the payment incurred by profit fees by shading their shouts to hide their actual trading surpluses, and then traders will prefer the marketplace charging a profit fee. Furthermore, we extended the analysis by considering multi-home and hybrid trading environments and different good properties. We then analysed the effects of different types of fees on obtaining market profits and keeping traders in a single marketplace environment, and showed that the transaction price percentage fee is the most effective in terms of making profits and at the same time keeping traders for the marketplace. Finally, we analysed how competing marketplaces set fees in equilibrium, and found that since traders can shade their shouts, the marketplace will charge a high profit fee.

This work addresses our research challenges of analysing equilibrium market selection and bidding strategies for traders and equilibrium charging strategies for marketplaces (i.e. research challenges 1, 2 and 4, see Section 1.2). Specifically, this is the first work that derives the equilibrium bidding strategies for traders in double auctions and analyses the effect of market fees on these strategies. This analysis is insightful and can be used to guide the design of a charging

	0.0	***0.1***	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
0.0	0.0,0.0	0.0,0.058	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.1	0.058,0.0	**0.190,0.190**	0.304,0.0	0.342,0.0	0.361,0.0	0.361,0.0	0.361,0.0	0.361,0.0	0.361,0.0	0.361,0.0	0.361,0.0
0.2	0.0,0.0	0.0,0.304	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.3	0.0,0.0	0.0,0.342	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.4	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.5	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.6	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.7	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.8	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
0.9	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0
1.0	0.0,0.0	0.0,0.361	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0	0.0,0.0

Table 4.1: Profits of marketplace 1 and marketplace 2. The first column is the registration fee of marketplace 1 and the first row is the registration fee of marketplace 2. The first element in each cell is marketplace 1's expected utility, and the second is marketplace 2's expected utility. Bold italic fees constitute a NEQ fee system.

	0.0	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	**0.9**	**1.0**
0.0	0.0,0.0	0.0,0.023	0.0,0.031	0.0,0.031	0.0,0.035	0.0,0.042	0.0,0.041	0.0,0.045	0.0,0.047	0.0,0.039	0.0,0.037
0.1	0.023,0.0	0.023,0.023	0.025,0.034	0.028,0.035	0.03,0.04	0.03,0.048	0.032,0.048	0.032,0.052	0.034,0.047	0.035,0.047	0.037,0.037
0.2	0.031,0.0	0.034,0.025	0.034,0.034	0.037,0.039	0.04,0.045	0.043,0.048	0.043,0.055	0.047,0.052	0.046,0.055	0.05,0.047	0.05,0.044
0.3	0.031,0.0	0.035,0.028	0.039,0.037	0.039,0.039	0.043,0.05	0.047,0.054	0.051,0.055	0.051,0.06	0.055,0.055	0.055,0.054	0.059,0.044
0.4	0.035,0.0	0.04,0.03	0.045,0.04	0.05,0.043	0.05,0.05	0.055,0.06	0.06,0.062	0.065,0.057	0.065,0.062	0.07,0.055	0.07,0.051
0.5	0.042,0.0	0.048,0.03	0.048,0.043	0.054,0.047	0.06,0.055	0.06,0.06	0.066,0.069	0.072,0.067	0.072,0.07	0.078,0.062	0.084,0.051
0.6	0.041,0.0	0.048,0.032	0.055,0.043	0.055,0.051	0.062,0.06	0.069,0.066	0.069,0.069	0.075,0.075	0.082,0.066	0.082,0.07	0.089,0.059
0.7	0.045,0.0	0.052,0.032	0.052,0.047	0.06,0.051	0.057,0.065	0.067,0.072	0.075,0.075	0.075,0.075	0.082,0.078	0.09,0.07	0.09,0.066
0.8	0.047,0.0	0.047,0.034	0.055,0.046	0.055,0.055	0.062,0.065	0.07,0.072	0.066,0.082	0.078,0.082	0.078,0.078	0.086,0.078	0.094,0.066
0.9	0.039,0.0	0.047,0.035	0.047,0.05	0.054,0.055	0.055,0.07	0.062,0.078	0.07,0.082	0.07,0.09	0.078,0.086	0.078,0.078	0.086,0.073
1.0	0.037,0.0	0.037,0.037	0.044,0.05	0.044,0.059	0.051,0.07	0.051,0.084	0.059,0.089	0.066,0.09	0.066,0.094	0.073,0.086	0.073,0.073

Table 4.2: Profits of marketplace 1 and marketplace 2. The first column is the profit fee of marketplace 1 and the first row is the profit fee of marketplace 2. The first element in each cell is marketplace 1's expected utility, and the second is marketplace 2's expected utility.

	0.0	**0.1**	**0.2**	**0.3**	**0.4**	**0.5**	**0.6**	**0.7**	**0.8**	***0.9***	**1.0**
0.0	0.0,0.0	0.0,0.023	0.0,0.029	0.0,0.031	0.0,0.035	0.0,0.042	0.0,0.041	0.0,0.045	0.0,0.047	0.0,0.039	0.0,0.037
0.1	0.058,0.0	0.058,0.041	0.058,0.055	0.058,0.07	0.058,0.09	0.058,0.107	0.078,0.116	0.097,0.119	0.078,0.133	**0.078,0.136**	0.097,0.125
0.2	0.0,0.0	0.0,0.046	0.0,0.06	0.0,0.078	0.0,0.101	0.0,0.12	0.0,0.136	0.0,0.149	0.0,0.157	0.0,0.159	0.0,0.153
0.3	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.157	0.0,0.149
0.4	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
0.5	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
0.6	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
0.7	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
0.8	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
0.9	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147
1.0	0.0,0.0	0.0,0.046	0.0,0.059	0.0,0.078	0.0,0.101	0.0,0.121	0.0,0.137	0.0,0.149	0.0,0.156	0.0,0.156	0.0,0.147

Table 4.3: Profits of marketplace 1 and marketplace 2. The first column is the registration fee of marketplace 1 and the first row is the profit fee of marketplace 2. The first element in each cell is marketplace 1's expected utility, and the second is marketplace 2's expected utility. Bold italic fees constitute a NEQ fee system.

strategy. For example, we found that when charging a profit fee, the marketplace may not be able to obtain a good level of profit even though a very high profit fee is charged, since traders will shade their shouts more. However, the marketplace can keep the number of traders at a good level. Another example is that charging a transaction fee is better than a registration fee to make profits and keep traders. In the following chapter of designing a competing marketplace, we will use these insights to design an effective charging policy.

Chapter 5

Designing a Competing Double Auction Marketplace

So far, we have analysed the equilibrium charging strategies for marketplaces and obtained a number of insights into individual and systemwide behaviour in a competing marketplace situation (see Chapters 3 and 4). However, as noted in Section 1.1, in addition to the design of a charging policy, we also need to design effective market policies that cover issues such as timing, matching, pricing and shout accepting policies. Now as we discussed in Section 2.5.1, there exist many such policies for competing marketplaces. However, we do not know which of them will perform well when marketplaces using different policies compete with each other. Therefore, we need to analyse how the different polices affect the performance of competing marketplaces. Specifically, in this chapter, we will conduct an experimental analysis on this issue in the context of the CAT competition since this provides an international benchmark for this problem. However, before we can do this, we first need to know how traders select marketplaces and submit shouts in this context since the traders' strategies will affect the effectiveness of the different market policies. Then, based on the analysis of traders' strategies, we can undertake an experimental analysis of market policies in different environments where different bidding strategies are adopted. By doing so, we obtain further insights about which policy is effective in guaranteeing traders' profits and thus attracting traders and how traders' bidding strategies affect the effectiveness of market policies. We then use these insights to design market policies for the CAT competition. This is in addition to the insights from Chapters 3 and 4 to designing a charging policy for the CAT competition.

The structure of this chapter is as follows. In Section 5.1, we describe how traders select marketplaces and submit shouts in the CAT competition. Then, in Section 5.2, we analyse how different market polices affect the performance of competing marketplaces when different bidding strategies are adopted. Through this analysis, we obtain several insights. In Section 5.3, we use the obtained insights to design market policies for the CAT competition. In Section 5.4, we design a charging policy and evaluate it in the CAT competition context. In Section 5.5, we describe the competition result of 2010 when our competing marketplace adopted the market policies and charging policy designed in this chapter. Finally, we summarise in Section 5.6.

5.1 Traders' Strategies in the CAT Competition

In this section, we describe how traders behave in the CAT competition (these are provided by the operators of the competition, not the competition entrants, see Cai et al. 2009; Niu et al. 2009). Firstly, we describe how traders select marketplaces. Then we analyse how traders submit shouts when using different bidding strategies.

5.1.1 Selecting Marketplaces

Firstly, we describe how traders select marketplaces in the CAT competition. In particular, the *ϵ-greedy exploration strategy* is adopted by all traders. This uses the parameter ϵ to adjust the balance between exploration of searching for the most profitable marketplace and exploitation of selecting what they believe is the most profitable marketplace (see Section 2.5.1). The parameter ϵ is from 0 to 1, and in the CAT competition, it is usually set as 0.1. In more detail, in the beginning of the competition, traders select each marketplace with equal probability. Then in the following days, according to their historical knowledge, they select the most profitable marketplace with probability $1 - \epsilon$, and explore other marketplaces with probability ϵ. Niu et al. (2007) experimentally analysed the traders' choices of marketplaces when they adopt this strategy in the specific context of the CAT competition. They showed that through exploration, traders will eventually locate the marketplace which is the most profitable. Therefore, since traders are able to search for the highly efficient marketplace when adopting this market selection strategy, competing marketplaces should provide effective market policies to maximise traders' profits to attract them. Note that this strategy is similar to our equilibrium analysis of traders' market selection strategies in Chapter 3, where we found that when we introduced randomisation for traders' market selection (i.e. traders randomly select marketplaces with some probability to explore other marketplaces), traders can explore to locate the cheapest marketplace (see Section 3.2.3.3).

5.1.2 Submitting Shouts

In the CAT competition, four heuristic bidding strategies are adopted: ZI-C, RE, GD and ZIP (see Section 2.3.2.1). In this section, we analyse how traders submit shouts when using these different bidding strategies. We first need to set up the marketplace where traders submit shouts. Specifically, we adopt the same market policies as those that were used in Chapters 3 and 4 (as we will show in Section 5.2, these market policies are highly efficient). In more detail, we adopt the round clearing policy (see Section 2.3.2.2), equilibrium matching policy (see Section

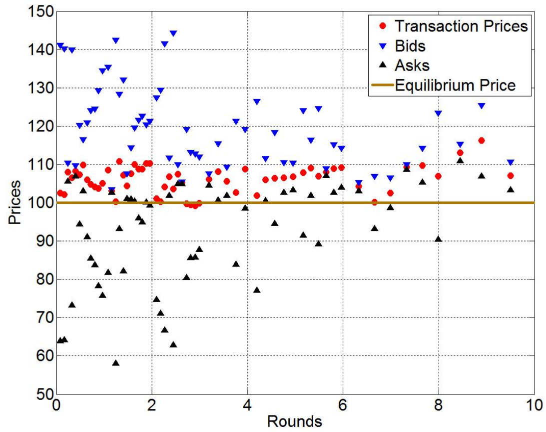

(a) ZI-C Strategy

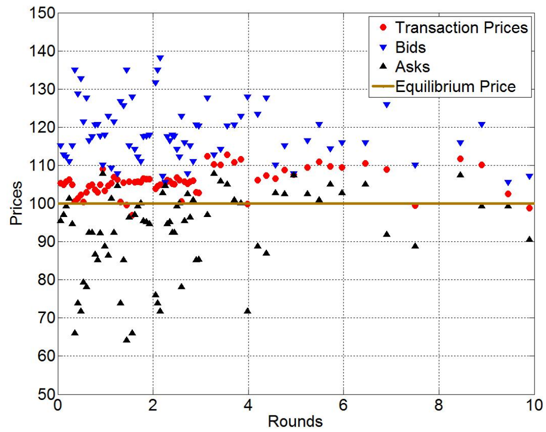

(b) RE Strategy

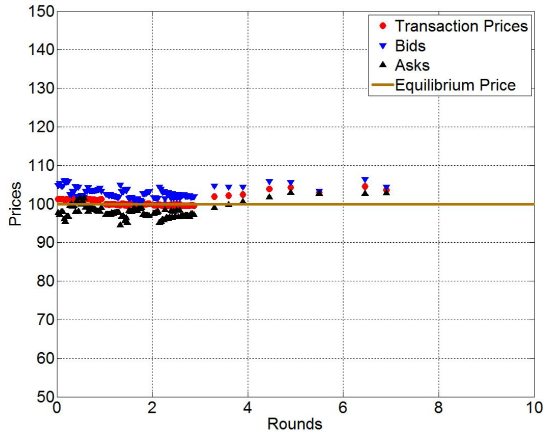

(c) ZIP Strategy

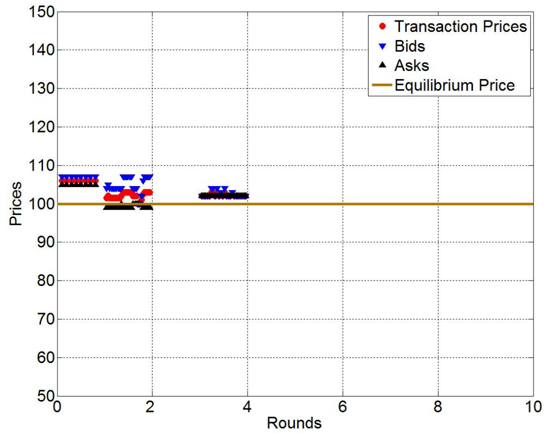

(d) GD Strategy

FIGURE 5.1: Bids, asks and transaction prices in the marketplace with traders using different bidding strategies.

2.3.2.2) and a k-pricing policy with $k = 0.5$ (see Section 2.3.2.2). For the shout accepting policy, we use the quote-beating accepting policy since this policy is widely used in real exchanges (like the NYSE, see also Section 2.3.2.2 for a detailed description of this policy). After setting up the marketplace, we now consider the composition of traders in the experiments. Specifically, we assume that there are 60 buyers and 60 sellers in this marketplace. Their private values are independently drawn from a uniform distribution between 50 and 150, and each trader is allowed to buy or sell up to 3 identical goods each day. Each experiment lasts for 100 days with 10 round per day and 1 second per round[1]. In the trading round, traders will submit shouts (and may improve their shouts in order to get matched). We run four experiments with all traders adopting ZI-C, RE, ZIP and GD strategies respectively based on the JCAT platform[2], and randomly choose one trading day for which we plot the bids, asks and transaction prices of all successful transactions in that day (the results are similar in other days).

[1]These setups are similar to those in the CAT competition.

[2]Here we use version 0.11 of JCAT platform. All the following experiments are also based on this version.

The experimental results are given in Figure 5.1 where we show the bids, asks and transaction prices of traders. From Figures 5.1(a) and 5.1(b), we can see that, although ZI-C and RE traders' transaction prices are close to the equilibrium price, which is 100, their shouts are typically far away from the equilibrium price (i.e. they shade their shouts very little). On the other hand, from Figures 5.1(c) and 5.1(d), we can see that the shouts from the ZIP and GD traders are concentrated in the area close to the equilibrium price (i.e. they shade their shouts a lot), and their transaction prices do indeed converge to the equilibrium price. The reason is as follows. Intra-marginal ZIP traders try to find shouts that provide high profits and also try to remain competitive in the marketplace. Through learning, ZIP traders find that, when their bids (asks) are a little higher (lower) than the equilibrium price, their profits are high and they are still sufficiently competitive to make transactions. Thus ZIP traders' shouts are concentrated in the area close to the equilibrium price. For traders using GD strategy, recall that in Section 2.3.2.1, we introduced the fact that the submitted shout depends on the belief of the shout accepted by the marketplace and the hidden profit (which is the difference between the private value and the shout). When their shouts are close to the equilibrium price, the hidden profits are high, and as shouts close to the equilibrium price are accepted by the marketplace, traders' beliefs about shout acceptance within the area close to the equilibrium price increase. Eventually, more traders will submit shouts close to the equilibrium price because of the high probability of acceptance and high hidden profits. We also note that in the last rounds, no transactions take place in the marketplace with trader adopting GD or ZIP strategy. This is because the more intelligent GD and ZIP traders can sell or buy items more quickly than ZI-C and RE traders. Moreover, we also run experiments by considering fees charged to traders. However, we find that fees cannot affect traders' shouts significantly when they use these heuristic bidding strategies, and the experimental results are similar to those in Figure 5.1. Furthermore, given this analysis, we further classify the four bidding strategies into two categories: the first is made of the ZI-C and RE strategies, where shouts generated by them are distributed over a wide range; the second consists of the ZIP and GD strategies, where their shouts are concentrated in the area close to the equilibrium price. This classification is useful because it enables us to group the four bidding strategies and can help us analyse the experiments of how each market policy affects the performance of competing marketplaces in the following section.

5.2 Experimental Analysis of Market Policies

In this section, we experimentally analyse how different market policies affect the performance of competing marketplaces in the context of the CAT competition. Specifically, we use allocative efficiency, market share, transaction success rate (TSR) and the number of transactions as metrics to measure marketplaces' performance. The reasons of adopting these metrics are as follows. Allocative efficiency is commonly used in literature to measure the performance of marketplaces; both market share and TSR are used as metrics in the CAT competition; and the number of transactions is important for us to determine whether to charge fees or not. This analysis will provide us insights for designing effective market polices in the CAT competition. Furthermore, recall that in Chapters 3 and 4, we conducted our equilibrium analysis by assuming that marketplaces adopt the specific market policies (i.e. round clearing, equilibrium matching and k-pricing policy with $k = 0.5$). Through the experimental analysis in this section, we will know whether these specific market policies perform well, and if so, then our assumption is reasonable. Note that since the effectiveness of market policies are affected by the traders' bidding strategies, we need to conduct the analysis in different environments with different bidding strategies. In the following, we first describe the different market policies that we will analyse. Then we run experiments in different environments to analyse how different market policies affect the marketplaces' performance.

5.2.1 Market Policies

In this section, we introduce the market policies that we want to analyse. For the timing policy, we consider continuous clearing (CC) and round clearing (RC) (see Section 2.3.2.2). For the matching policy, we consider the equilibrium matching policy (ME) which matches buyers with high bids with sellers with low asks, i.e. match intra-marginal buyers with intra-marginal sellers, and the maximising volume matching policy (MV)[3] which matches buyers with high bids with sellers with high asks if they can be matched, i.e. match intra-marginal buyers with extra-marginal sellers (see Section 2.3.2.2). For the pricing policy, existing research has shown that the k-pricing policy with k=0.5 is efficient both in terms of the allocative efficiency and the traders' efficiency (see Section 2.3.2.2), and thus we use this pricing policy in our analysis. In terms of the shout accepting policy, we mainly consider the two alternatives that have been most widely used in the CAT Competition: the *quote-beating accepting policy* (AQ) and the estimated equilibrium accepting policy (AE) (see Section 2.3.2.2). Note that in addition to these market policies, other policies do exist. However, from the CAT competition, we find that lots of policies designed by entrants are built from these specific policies. Thus here we focus on these market policies.

In addition to the above matching policies, we also design a specific one. As discussed in Section 2.3.2.2, there are some disadvantages with both ME and MV. The former aims to maximise traders' profits. Thus, intra-marginal traders are happy with this matching policy since they can usually earn high expected profits. However, it is difficult for marginal traders, who submit shouts close to the equilibrium price, to make transactions. The latter aims to maximise the number of transactions, and thus benefits TSR. In this case, marginal traders are easily matched, and thus they favour this policy. However, using this policy, intra-marginal traders cannot be guaranteed to trade with other intra-marginal traders to obtain high expected profit. Thus their profit is lost and it causes a decrease in allocative efficiency. Thus in order to strike a balance between maximising profits and maximising the number of transactions, we developed the following matching policy which combines the advantages of both ME and MV. At the end of each

[3]We had to implement the MV matching policy since it was not actually implemented in the JCAT platform.

Marketplace Policies	Timing Policy	Matching Policy	Shout Accepting Policy
CC-ME-AQ	continuous clearing	equilibrium matching	quote-beating accepting
CC-MV-AQ	continuous clearing	maximising volume	quote-beating accepting
RC-ME-AQ	round clearing	equilibrium matching	quote-beating accepting
RC-MV-AQ	round clearing	maximising volume	quote-beating accepting
RC-MEV-AQ	round clearing	combined matching of ME and MV	quote-beating accepting
CC-ME-AE	continuous clearing	equilibrium matching	estimated equilibrium accepting
CC-MV-AE	continuous clearing	maximising volume	estimated equilibrium accepting
RC-ME-AE	round clearing	equilibrium matching	estimated equilibrium accepting
RC-MV-AE	round clearing	maximising volume	estimated equilibrium accepting
RC-MEV-AE	round clearing	combined matching of ME and MV	estimated equilibrium accepting

TABLE 5.1: The different marketplace policies in the experimental setup.

round, the marketplace is firstly cleared using ME, which guarantees that intra-marginal traders obtain high expected profits. Then, when the bid-ask spread[4] is less than a certain threshold, i.e. most remaining traders are marginal traders with low expected profits, the marketplace is cleared using MV. By so doing, it generates high profits for most traders, similar to ME, and increases the number of transaction, similar to MV. In the following, this matching policy is referred to as MEV. Note that we only adopt the MEV matching policy when the marketplace is cleared at the end of each round, i.e. in combination with the RC policy. When the marketplace is cleared continuously, both ME and MV cannot guarantee high profits of intra-marginal traders. Thus MEV cannot offer any advantage in this case, and we will not combine CC with MEV in our experiments.

From the above we therefore consider two types of timing policies, three types of matching policies, two types of shout accepting policies and one type of pricing policy. Given this, we consider 10 different marketplaces with different combinations of policies, namely: CC-ME-AQ, CC-ME-AE, CC-MV-AQ, CC-MV-AE, RC-ME-AQ, RC-ME-AE, RC-MEV-AQ, RC-MEV-AE, RC-MV-AQ, RC-MV-AE. See Table 5.1 for an overview of these policies. In what follows, we analyse the above policies in the competing marketplace context with different bidding strategies. We also want to know whether the marketplace adopting RC-ME (i.e. the market policies we used in Chapters 3 and 4) can perform well when competing with marketplaces using other combinations.

5.2.2 Experiments with Homogeneous Bidding Strategy

We first run experiments in the environments where all traders use the same bidding strategy (i.e. homogeneous bidding strategy). From this, we will know how each bidding strategy affects the effectiveness of market policies, and this will provide the foundation for understanding the further analysis where different traders may use different bidding strategies (i.e. heterogeneous bidding strategies). The experimental setup is as follows. Each experiment runs for 100 days with 10 rounds per day and 1 second per round. There are 200 buyers and 200 sellers. The private values of all traders are independently drawn from a uniform distribution between 50 and

[4]This is the difference between the outstanding bid and the outstanding ask in the marketplace, see Section 2.3.2.1 for the definition of outstanding bid/ask.

150, and each trader is allowed to buy or sell up to 3 goods per day. All traders use either ZI-C, RE, GD or ZIP strategy. Furthermore, in the experiments, there are 10 competing marketplaces and each one of them adopts one of the policy combinations specified in Table 5.1. In total, we run 4 experiments with all traders adopting ZI-C, RE, GD and ZIP strategies respectively, and each experiment is repeated 40 times.

The results with all traders only adopting ZI-C strategy are shown in Figure 5.2 and Table 5.2. For clarity, we do not add error bars in the figures, and instead we calculate 95% confidence intervals of the average values of each metric in tables. The average values are taken over 100 days and 40 runs. The confidence intervals show the error of the average daily result over the 40 runs. For example, in Table 5.2, 91.171±1.211 means the 95% confidence interval range of allocative efficiency of the marketplace adopting CC-ME-AE, where 91.171 − 1.211 is the lower 95% confidence limit, 91.171 + 1.211 is the upper 95% confidence limit and 91.171 is the mean of scores of 100 days and 40 runs. Here we can see that the marketplaces using RC-ME-AQ and RC-MEV-AQ, significantly outperform other marketplaces in terms of allocative efficiency (see Figure 5.2(a) and Table 5.2), and also in terms of market share (see Figure 5.2(b) and Table 5.2) and the number of transactions completed (see Figure 5.2(d) and Table 5.2). The reason for this is that ZI-C traders submit shouts from a uniform distribution, and, as a result, their shouts are not close to the equilibrium price (as shown in Figure 5.1(a)). Consequently, when the marketplace maximises the number of transactions, as the marketplaces using CC-ME, CC-MV and RC-MV do, the intra-marginal buyers cannot be guaranteed to trade with the intra-marginal sellers, and thus lose potential profits. These intra-marginal traders will then choose other more profitable marketplaces. Furthermore, AE provides a tighter restriction on accepting shouts, and thus, the TSR of marketplaces adopting AE is significantly better than a marketplace using AQ (see Figure 5.2(c) and Table 5.2). However, because of the imprecise estimation of the equilibrium price, this accepting policy drives some traders to leave the marketplace, which decreases slightly the number of transactions. From this experiment, we find that the marketplace with RC-ME-AQ performs well in all these metrics. In addition, from Figure 5.2(b), we note that in the first few trading days, the market shares of marketplaces using CC-ME-AQ, CC-MV-AQ and RC-MV-AQ are high. This is because some extra-marginal traders may be able to “steal” transactions, and thus prefer these marketplaces. However, because intra-marginal traders’ profits are harmed, they leave these marketplaces and this leads to a decreased market share.

The second experiment uses RE traders. These experimental results are shown in Figure 5.3 and Table 5.3. Now, recall from Section 5.1.2 that shouts of RE traders are also far away from the equilibrium price, and therefore, as we expect, the experimental results are similar to those experiments with ZI-C traders and the explanation is the same as above.

Figures 5.4 and 5.5 and Tables 5.4 and 5.5 show the experimental results with ZIP and GD traders respectively. As before, since shouts of both ZIP and GD traders are close to the equilibrium price, these two sets of experiments show the same pattern. From this, we can see that the allocative efficiency of marketplaces adopting AE is quite low (see Figures 5.4(a) and

5.5(a), and Tables 5.4 and 5.5), and the number of transactions is very low (see Figures 5.4(d) and 5.5(d) and Tables 5.4 and 5.5). The reasons for this are as follows. As discussed in Section 5.1.2, ZIP and GD traders' shouts are close to the equilibrium price, and so transaction prices converge to the equilibrium price. However, the estimated equilibrium price is not always precise enough. A small error of estimated equilibrium price may cause the rejection of many shouts, and then traders will choose to leave the marketplace[5]. Therefore, a marketplace using AE performs worse than a marketplace using AQ. From the experimental results, we can also determine that the market share in the marketplaces adopting CC-ME-AQ or RC-MV-AQ is slightly higher than the marketplaces using RC-ME-AQ (see Tables 5.4 and 5.5, the result is statistically significant). This is different from the experiments with the ZI-C and RE traders. The reason for this discrepancy is that ZIP traders' shouts are close to the equilibrium prices, and thus the marketplace cannot distinguish between intra-marginal and marginal traders. Therefore, the advantage of RC-ME in maximising traders' profits is weakened. As a result, in this case, the performance of RC-ME is similar to RC-MV and CC-ME in terms of the traders' profits. However, in marketplaces using RC-MV and CC-ME, marginal traders can easily be matched, and some extra-marginal traders may also trade. Therefore, marginal traders and some extra-marginal traders prefer marketplaces using RC-MV and CC-ME. However, because the number of marginal traders and extra-marginal traders that steal transactions is small, and they are distributed in marketplaces using RC-MV and CC-ME respectively, CC-ME-AQ and RC-MV-AQ just slightly outperform RC-ME-AQ in terms of market share.

Now we have run experiments where all traders use the same bidding strategy. In summary, because of the similarity of traders using ZI-C and RE strategies, the experimental results using these traders are similar. Similarly, we find that experimental results with ZIP traders are similar to those with GD traders. When traders' shouts are far away from the equilibrium price, the timing and matching policies mainly determine the marketplace's performance. Conversely, when traders' shouts are close to the equilibrium price, the shout accepting policy mainly determines the marketplace's performance. Specifically, from our experimental results, we can see that the performance of marketplaces using AE in the environments with GD and ZIP strategies, and marketplaces using CC-ME, CC-MV, RC-MV in the environments with ZI-C and the RE strategies is very poor. We also find that the marketplace using RC-ME-AQ performs well in all cases, even though they might be outperformed slightly by other marketplaces in some particular environments.

5.2.3 Experiments with Heterogenous Bidding Strategies

In the above, we analysed how different market policies affect the competing marketplaces' performance in the environments where all traders use the homogeneous bidding strategy. How-

[5]In the above experiments with ZI-C or RE traders, because traders' shouts are generally far away from the equilibrium price, the error of the estimated equilibrium price has little impact on the number of rejected shouts in a marketplace using AE. This is why the AE has small impact in such a marketplace, in contrast to the experiments with ZIP traders.

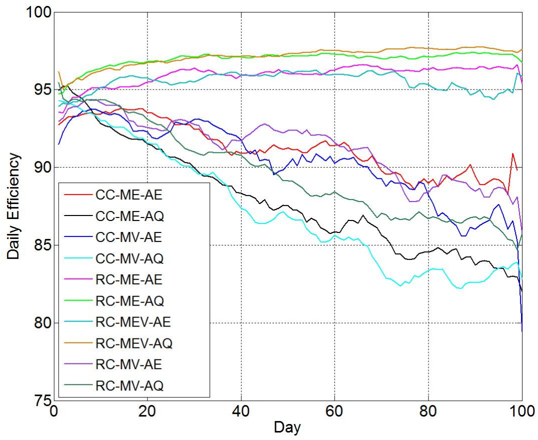

(a) Allocative efficiency

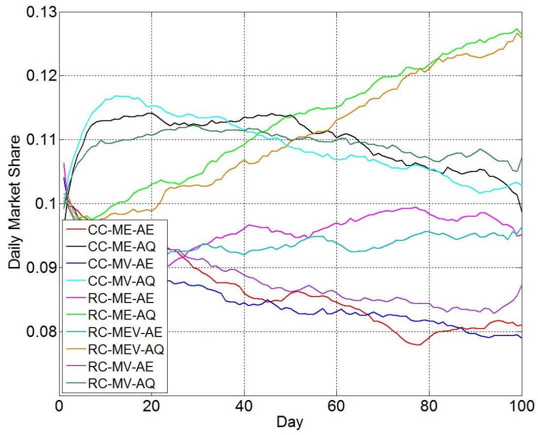

(b) Market share

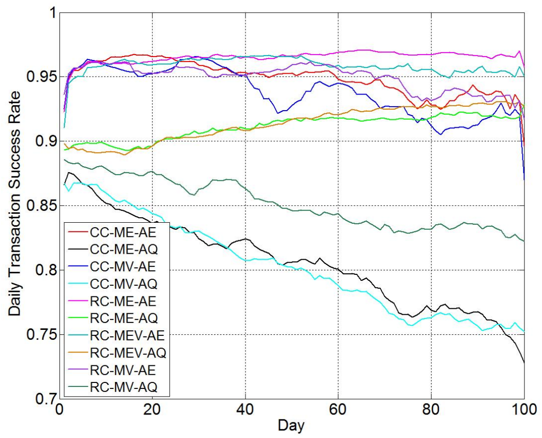

(c) TSR

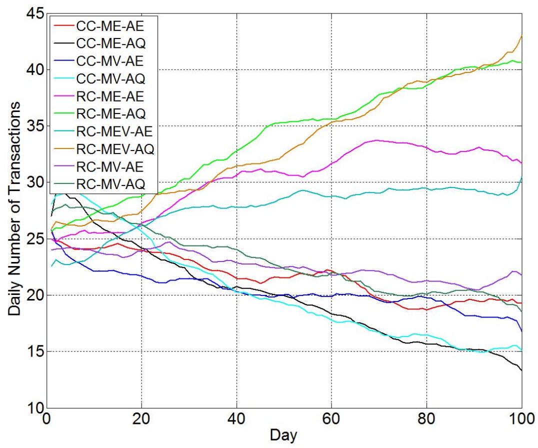

(d) Number of transactions

FIGURE 5.2: Scores of marketplaces with ZI-C strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	91.171±1.211	0.087±0.003	0.949±0.014	21.638±1.923
CC-ME-AQ	87.900±0.753	0.110±0.004	0.806±0.016	20.018±1.578
CC-MV-AE	90.140±2.002	0.085±0.004	0.939±0.020	20.382±1.956
CC-MV-AQ	87.193±0.856	0.109±0.005	0.803±0.015	20.342±2.122
RC-ME-AE	95.951±0.360	0.096±0.005	**0.965±0.003**	30.228±2.886
RC-ME-AQ	**96.970±0.335**	**0.113±0.006**	0.911±0.012	**34.089±3.005**
RC-MEV-AE	95.521±0.665	0.093±0.005	**0.959±0.007**	27.685±2.818
RC-MEV-AQ	**97.118±0.175**	**0.110±0.005**	0.913±0.007	**33.312±2.422**
RC-MV-AE	91.324±1.650	0.088±0.003	0.950±0.017	22.579±1.837
RC-MV-AQ	89.623±0.730	0.109±0.004	0.852±0.013	23.007±1.950

TABLE 5.2: Average daily results and corresponding 95% confidence intervals of marketplaces with ZI-C strategy. The average values are taken over 100 days and 40 runs. The confidence intervals show the error of the average daily result over the 40 runs. Bold face indicates that the corresponding marketplaces have high value ranges of the metric.

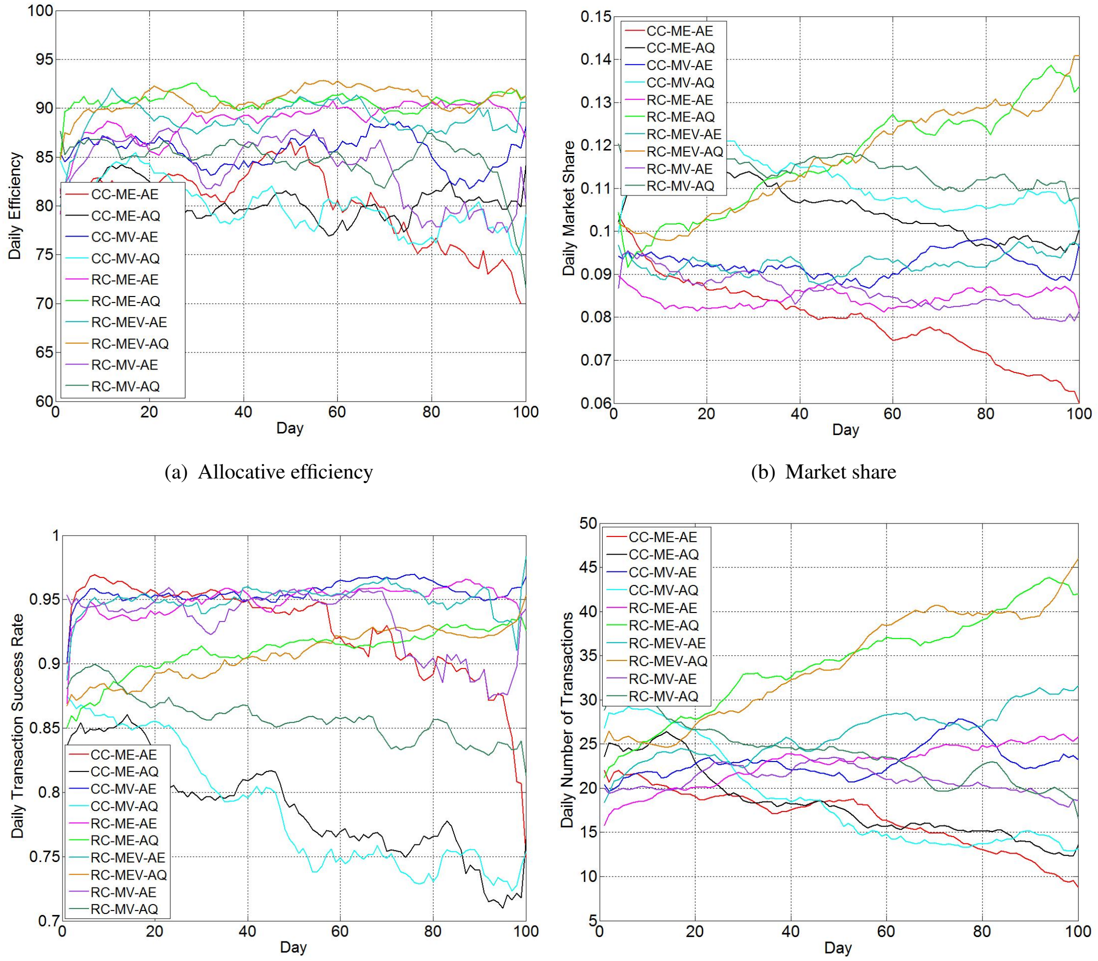

(a) Allocative efficiency

(b) Market share

(c) TSR

(d) Number of transactions

FIGURE 5.3: Scores of marketplaces with RE strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	80.024±1.532	0.080±0.003	0.928±0.012	16.632±1.608
CC-ME-AQ	80.456±0.663	0.106±0.004	0.788±0.010	18.126±1.178
CC-MV-AE	85.573±0.616	0.092±0.004	0.956±0.003	22.994±1.805
CC-MV-AQ	80.008±0.859	0.112±0.003	0.786±0.011	18.832±1.021
RC-ME-AE	88.88±0.332	0.084±0.004	**0.950±0.004**	22.714±1.956
RC-ME-AQ	**90.770±0.237**	**0.117±0.004**	0.910±0.007	**34.228±2.429**
RC-MEV-AE	88.858±0.670	0.092±0.003	**0.950±0.009**	26.124±2.372
RC-MEV-AQ	**90.910±0.296**	**0.117±0.004**	0.907±0.008	**34.122±2.838**
RC-MV-AE	83.743±1.595	0.086±0.004	0.933±0.015	20.92±2.264
RC-MV-AQ	84.804±0.565	0.115±0.002	0.858±0.004	24.098±0.827

TABLE 5.3: Average daily results and corresponding 95% confidence intervals of marketplaces with RE strategy.

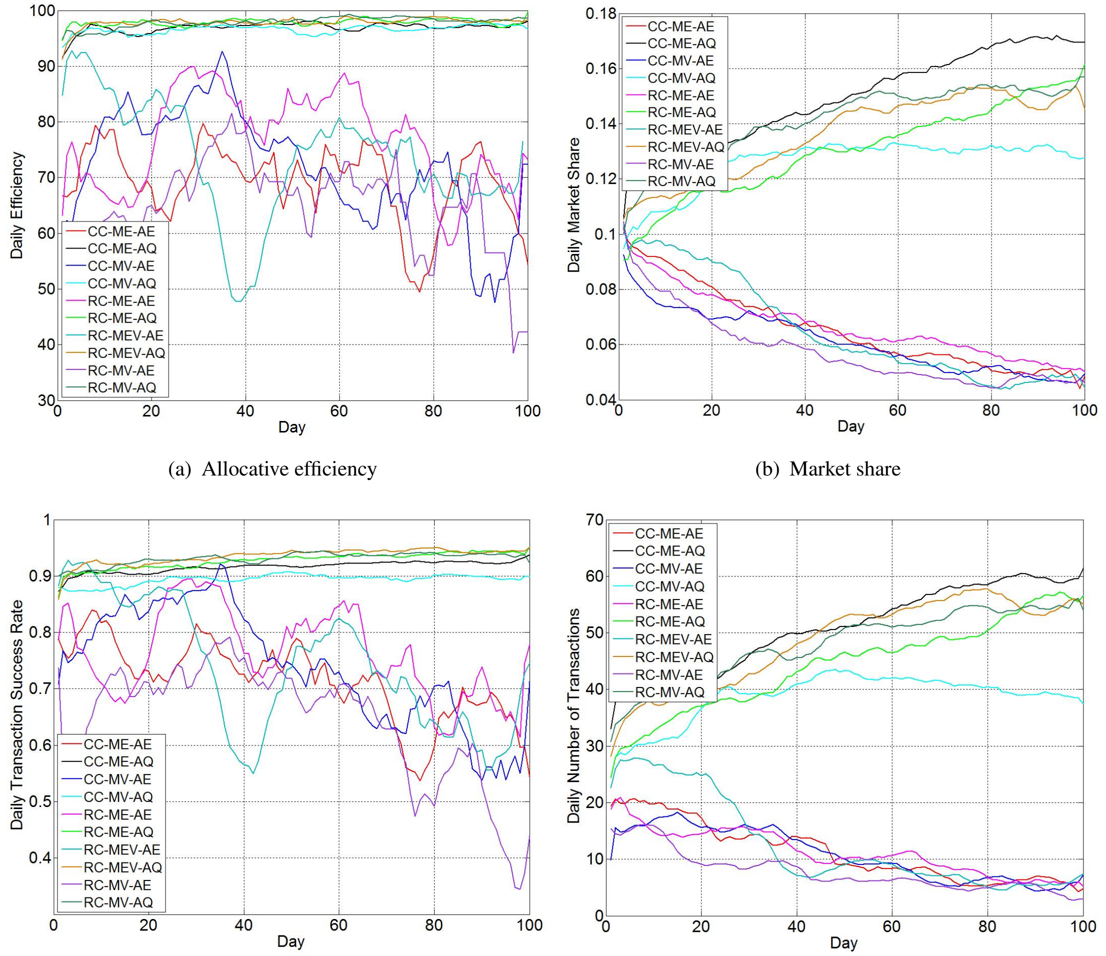

(a) Allocative efficiency (b) Market share

(c) TSR (d) Number of transactions

FIGURE 5.4: Scores of marketplaces with ZIP strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	67.538±5.030	0.066±0.003	0.715±0.050	11.286±1.830
CC-ME-AQ	96.952±0.108	**0.149±0.003**	0.917±0.001	**50.758±1.275**
CC-MV-AE	70.935±4.663	0.061±0.003	0.738±0.047	10.8±1.621
CC-MV-AQ	96.627±0.196	0.126±0.005	0.894±0.006	38.666±2.802
RC-ME-AE	75.14±2.597	0.067±0.002	0.761±0.025	11.192±0.683
RC-ME-AQ	**97.873±0.058**	0.130±0.004	**0.928±0.003**	44.352±2.033
RC-MEV-AE	72.224±2.814	0.066±0.002	0.741±0.028	12.616±1.146
RC-MEV-AQ	**97.897±0.105**	0.136±0.003	**0.936±0.002**	**48.57±1.86**
RC-MV-AE	61.998±5.574	0.058±0.004	0.639±0.052	7.806±1.733
RC-MV-AQ	**97.882±0.123**	**0.142±0.004**	**0.931±0.003**	**47.992±1.633**

TABLE 5.4: Average daily results and corresponding 95% confidence intervals of marketplaces with ZIP strategy.

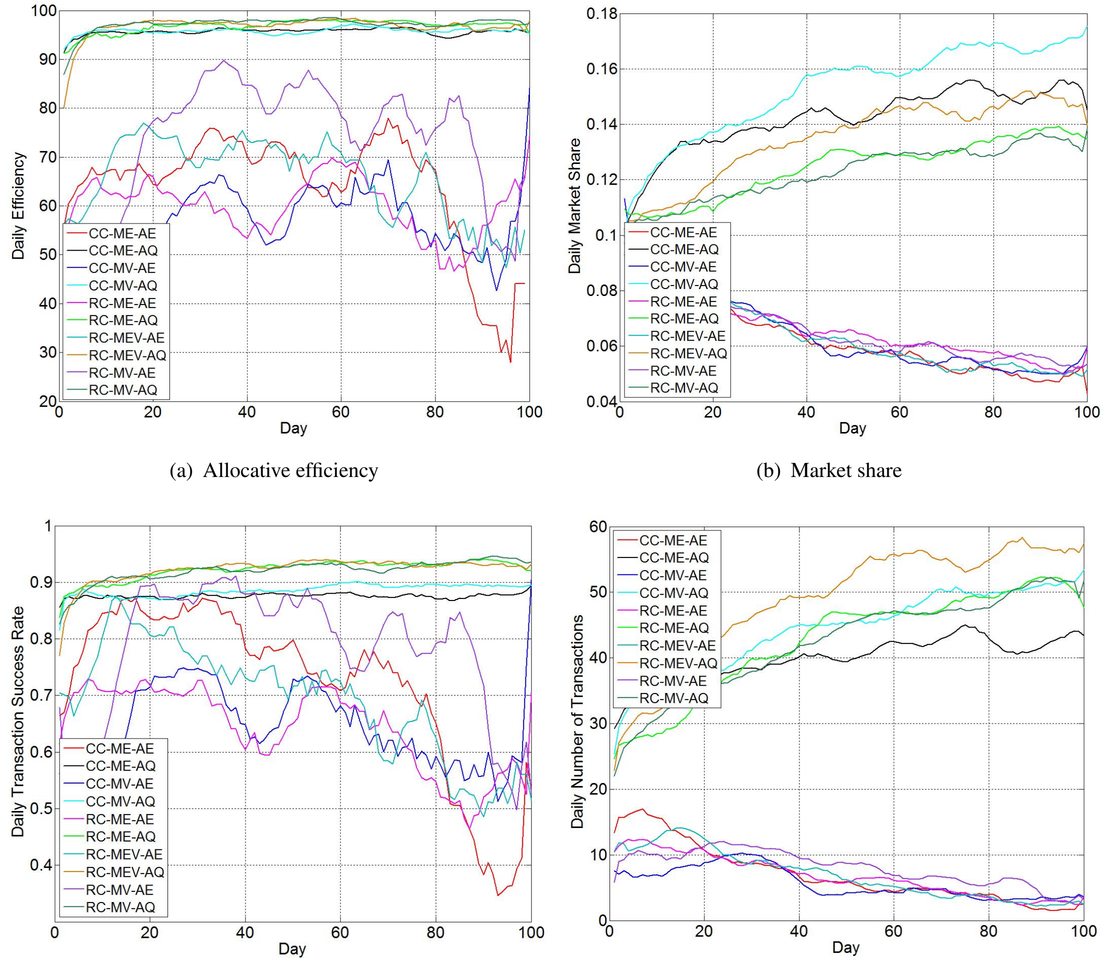

(a) Allocative efficiency

(b) Market share

(c) TSR

(d) Number of transactions

FIGURE 5.5: Scores of marketplaces with GD strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	63.248±0.860	0.064±0.001	0.717±0.016	7.03±0.407
CC-ME-AQ	95.686±0.315	**0.142±0.004**	0.876±0.007	39.966±2.146
CC-MV-AE	53.789±4.835	0.064±0.002	0.633±0.058	5.676±0.834
CC-MV-AQ	95.834±0.170	**0.153±0.004**	0.886±0.003	**44.518±1.531**
RC-ME-AE	59.056±9.185	0.066±0.003	0.642±0.099	6.772±1.168
RC-ME-AQ	**96.830±0.096**	0.125±0.002	**0.923±0.001**	**42.64±0.482**
RC-MEV-AE	63.216±4.352	0.064±0.002	0.692±0.040	6.94±0.797
RC-MEV-AQ	**96.808±0.156**	**0.135±0.003**	**0.923±0.004**	**48.662±2.327**
RC-MV-AE	72.220±2.261	0.065±0.001	0.778±0.024	8.07±0.686
RC-MV-AQ	**97.272±0.112**	0.123±0.003	**0.921±0.003**	**42.452±1.953**

TABLE 5.5: Average daily results and corresponding 95% confidence intervals of marketplaces with GD strategy.

ever, it often happens that traders use heterogeneous bidding strategies in the CAT competition. Therefore, in this section, we analyse how different market policies affect the performance of marketplaces in the environments where traders use heterogeneous bidding strategies. Specifically, we will analyse different environments where a big proportion of traders adopts a certain bidding strategy. We also want to analyse the environment where the same number of traders adopts each bidding strategy. In more detail, we consider five different environments where traders adopting heterogenous bidding strategies: (i) 300 ZI-C, 40 RE, 40 GD and 40 ZIP traders; (ii) 40 ZI-C, 300 RE, 40 GD and 40 ZIP traders; (iii) 40 ZI-C, 40 RE, 300 GD and 40 ZIP traders; (iv) 40 ZI-C, 40 RE, 40 GD and 300 ZIP traders; (v) 100 ZI-C, 100 RE, 100 GD and 100 ZIP traders. In all of these environments, the number of buyers using a particular bidding strategy is equal to the number of sellers[6].

The result when using environment (i) are shown in Figure 5.6 and Table 5.6. From these experiments, we can see that the marketplace adopting RC-ME-AQ still performs well (see allocative efficiency in Figure 5.6(a) and Table 5.6 and the number of transactions in Figure 5.6(d) and Table 5.6). We also note that the performance of marketplaces adopting RC-ME-AE improves compared with the experimental results with traders adopting the ZI-C strategy (see Figure 5.2 and Table 5.2). The reason is as follows. Shouts from ZIP and GD traders in this environment generate transaction prices closer to the equilibrium price than those with traders adopting ZI-C strategy. This means that the marketplace using AE can estimate the equilibrium price more precisely. Because shouts of ZI-C traders are far away from the equilibrium price, as previously discussed, the estimation error of the equilibrium price has little impact on the number of rejected shouts. Thus the marketplace adopting RC-ME-AE improves within this environment. Then we run experiments with environment (ii) where most traders use the RE strategy, we can see the experimental results shown in Figure 5.7 and Table 5.7 are similar to those with environment (i). The marketplace adopting RC-ME-AQ still performs well.

Our third set of experiments uses environment (iii) and the results are shown in Figure 5.8 and Table 5.8. We can see that the marketplace adopting RC-ME-AQ still performs well. Furthermore, similar to previous experiments with all traders adopting the GD strategy, we can see that the marketplaces using AE performs badly because of the rejection of many shouts. We also find that compared with the experimental results with trader only adopting GD strategy (see Figure 5.5 and Table 5.5), the marketplaces adopting AE improve their performance slightly in this environment with heterogenous strategies. The main reason for this is that in this environment, the beliefs of GD traders on shouts which are a bit further away from the equilibrium price increase because of the wide range of ZI-C and RE traders' shouts (see the calculation of belief function in equation 2.8 and 2.9 of Section 2.3.2.1), and thus the shouts of the GD traders in this environment become a little wider than those of all traders adopting the GD strategy only. This means that the number of rejected shouts when adopting AE decreases, and thus the performance of the marketplaces using AE improves slightly. It also explains why in this environment the marketplaces adopting CC-ME, CC-MV and RC-MV perform slightly worse than they do

[6]This is the same as that of the CAT competition in recent years.

with all traders adopting the GD strategy, since now, intra-marginal traders are more likely to be matched with marginal traders, and thus choose to leave these marketplaces. Furthermore, the experimental results with environment (iv) where most traders use ZIP strategy are shown in Figure 5.9 and Table 5.9. We find that they are similar to those with the environment (iii) where most traders use GD strategy.

Finally, the experimental results with environment (v) are shown in Figure 5.10 and Table 5.10. In this environment, the numbers of traders using each of the ZI-C, RE, ZIP and GD strategies are equal. Again we can see that the marketplace using RC-ME-AQ performs well in the allocative efficiency (see Figure 5.10(a)) and in the number of transactions (see Figure 5.10(d)).

Now we have analysed how different market policies affect the allocative efficiency, market share, TSR and the number of transactions in different environments. This analysis provides us insights of designing market policies for the CAT competition. In the following section, we will describe our design of market policies in detail.

5.3 Market Policy Design

In this section, we use insights from the above analysis to design market policies for the CAT competition. In the experimental analysis, we found that the marketplace adopting RC-ME-AQ always performs well in different environments. Therefore, our design of market policies will be built from RC-ME-AQ. In the following, we discuss how to design each market policy in detail.

For the timing policy, in Section 5.2, we found that when traders use the GD or the ZIP strategy, the continuous clearing policy also performs well. Actually, from Section 2.3.2.1, we can see that GD traders need to use information about transactions to improve their beliefs of shout acceptance. Therefore, in our timing policy, in the first round of each trading day, when a new shout is accepted by the marketplace, and this shout is close to the equilibrium price, the marketplace will immediately clear the shout (if it can be matched). However, if this shout is far away from the equilibrium price, it will be cleared at the end of the round. By so doing, on one hand, we can guarantee profits for buyers with high limit prices and sellers with low cost prices (since these traders are more likely to be matched at the end of round); on the other hand, the immediately matched bids and asks will provide useful information for traders to improve their shouts. After the first round, traders will have information to improve their shouts. Then in the remaining rounds, our marketplace will use round clearing policy in order to guarantee traders' profits.

For the matching policy, as we discussed in Section 5.2, the ME matching policy can maximise traders' profits, and thus it can attract traders to improve the market share. However, it is difficult for marginal traders to make transactions. The MV matching can maximise the number of tran-

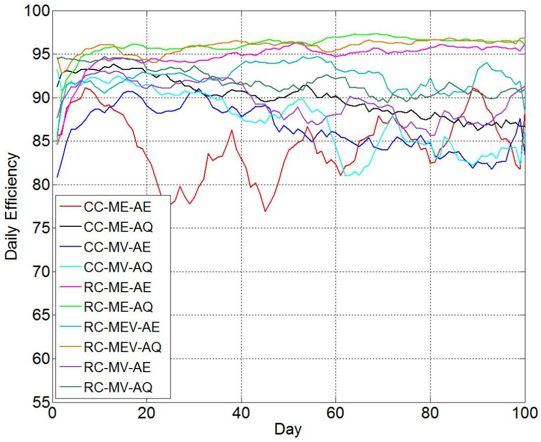

(a) Allocative efficiency

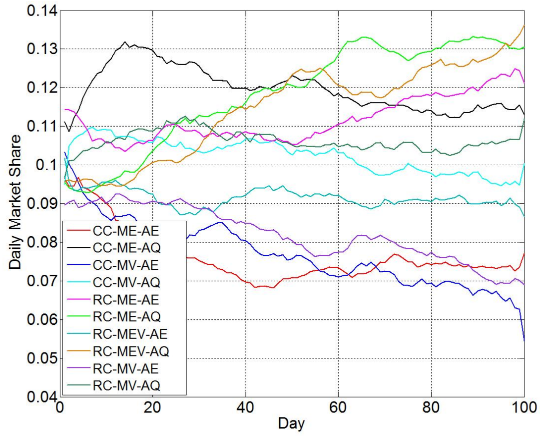

(b) Market share

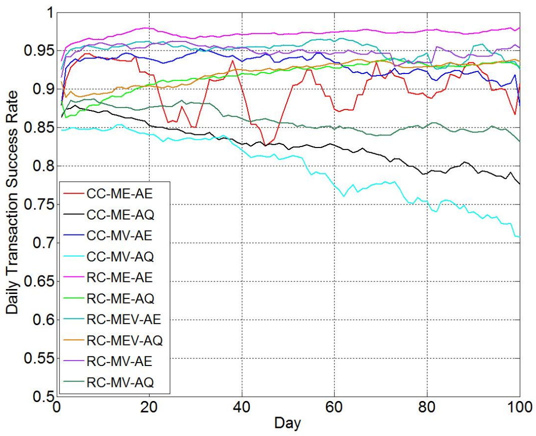

(c) TSR

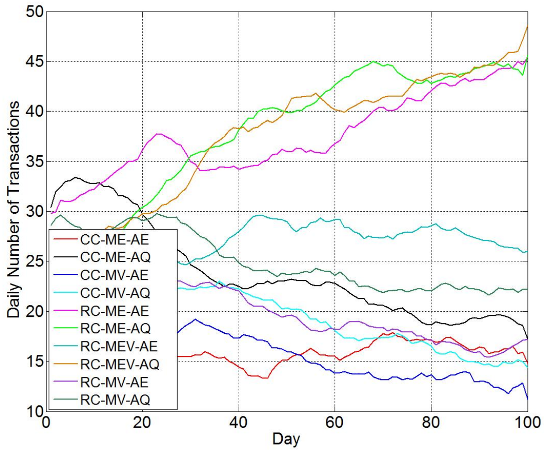

(d) Number of transactions

FIGURE 5.6: Scores of marketplaces in the environment where most traders use ZI-C strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	84.582±3.179	0.076±0.003	0.902±0.031	16.742±1.849
CC-ME-AQ	90.069±0.725	**0.120±0.003**	0.8280±0.008	23.938±0.620
CC-MV-AE	86.438±0.521	0.077±0.002	0.932±0.005	15.78±1.030
CC-MV-AQ	87.489±1.212	0.103±0.005	0.798±0.018	20.268±1.977
RC-ME-AE	94.730±0.355	0.112±0.002	**0.972±0.002**	**37.504±1.231**
RC-ME-AQ	**96.067±0.101**	**0.118±0.003**	0.918±0.006	**37.842±2.689**
RC-MEV-AE	92.423±1.097	0.092±0.005	**0.951±0.009**	26.962±2.883
RC-MEV-AQ	**95.886±0.317**	0.115±0.004	0.921±0.005	**37.572±2.398**
RC-MV-AE	89.655±0.514	0.082±0.002	**0.950±0.003**	20.054±0.599
RC-MV-AQ	91.945±0.469	0.106±0.002	0.861±0.010	25.018±1.543

TABLE 5.6: Average daily results and corresponding 95% confidence intervals of marketplaces in the environment where most traders use ZI-C strategy.

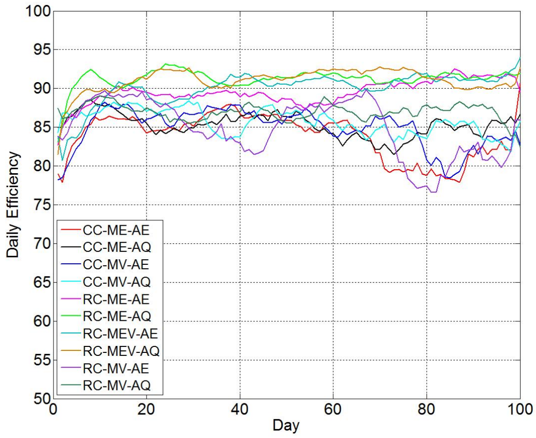

(a) Allocative efficiency

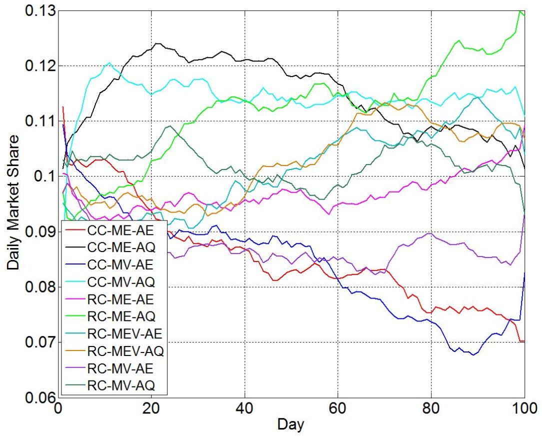

(b) Market share

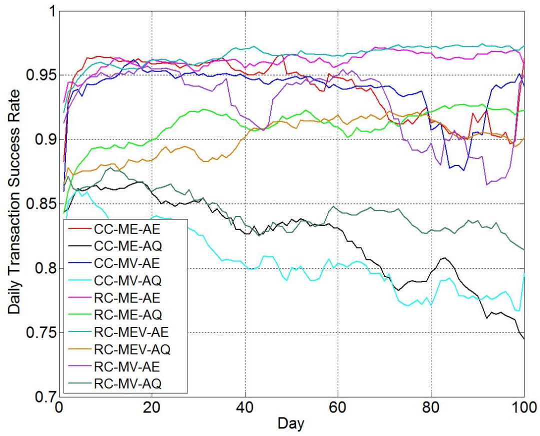

(c) TSR

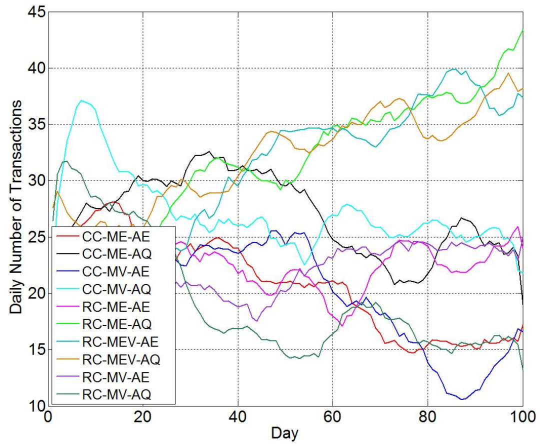

(d) Number of transactions

FIGURE 5.7: Scores of marketplaces in the environment where most traders use RE strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	83.809±1.227	0.086±0.003	0.940±0.010	20.69±1.649
CC-ME-AQ	85.304±1.344	**0.115±0.003**	0.824±0.013	25.972±1.853
CC-MV-AE	84.906±1.144	0.084±0.002	0.939±0.008	20.136±1.695
CC-MV-AQ	85.695±1.784	**0.115±0.004**	0.807±0.018	25.218±2.199
RC-ME-AE	89.560±0.259	0.097±0.003	**0.961±0.003**	20.584±2.228
RC-ME-AQ	**91.334±0.247**	0.112±0.003	0.911±0.004	**31.104±1.461**
RC-MEV-AE	90.007±0.562	0.101±0.002	**0.965±0.001**	**30.976±1.607**
RC-MEV-AQ	**91.071±0.421**	0.103±0.004	0.901±0.005	**31.316±1.496**
RC-MV-AE	84.845±2.084	0.087±0.005	0.928±0.016	21.264±3.382
RC-MV-AQ	86.962±0.881	0.102±0.002	0.844±0.006	19.018±0.922

TABLE 5.7: Average daily results and corresponding 95% confidence intervals of marketplaces in the environment where most traders use RE strategy.

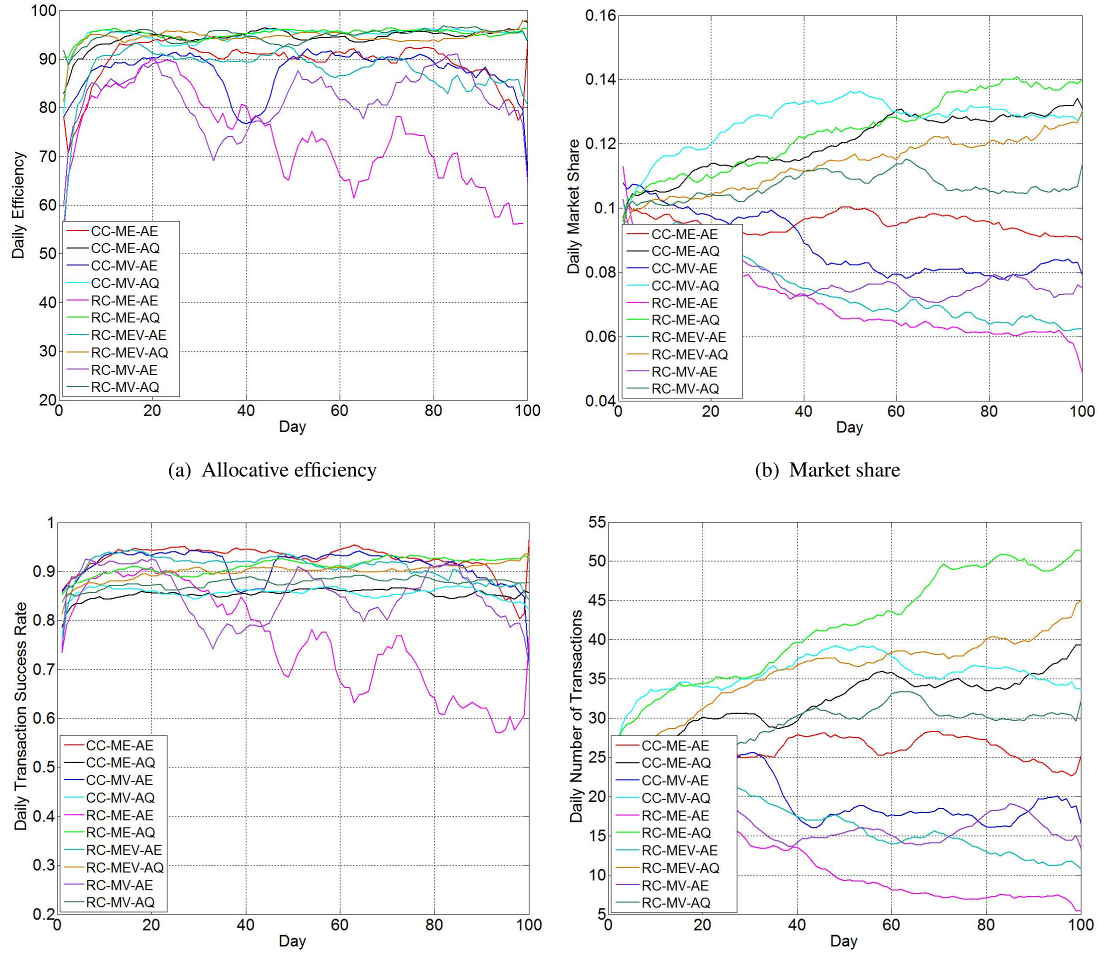

(a) Allocative efficiency

(b) Market share

(c) TSR

(d) Number of transactions

FIGURE 5.8: Scores of marketplaces in the environment where most traders use GD strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	89.279±1.498	0.095±0.007	**0.924±0.011**	25.488±4.013
CC-ME-AQ	**94.504±0.472**	**0.120±0.005**	0.854±0.008	32.176±2.538
CC-MV-AE	87.649±1.192	0.089±0.004	**0.914±0.008**	20.472±1.786
CC-MV-AQ	**95.002±0.542**	**0.127±0.007**	0.857±0.014	**35.552±4.048**
RC-ME-AE	72.519±4.413	0.070±0.005	0.758±0.042	11.204±2.344
RC-ME-AQ	**95.261±0.291**	**0.124±0.004**	**0.912±0.007**	**41.518±2.669**
RC-MEV-AE	88.064±0.887	0.076±0.002	**0.910±0.008**	17.126±1.125
RC-MEV-AQ	**94.589±0.266**	0.114±0.004	0.900±0.005	**35.866±2.386**
RC-MV-AE	81.88±3.610	0.079±0.005	0.853±0.035	16.606±2.700
RC-MV-AQ	**94.863±0.456**	0.107±0.004	0.877±0.006	29.396±2.045

TABLE 5.8: Average daily results and corresponding 95% confidence intervals of marketplaces in the environment where most traders use GD strategy.

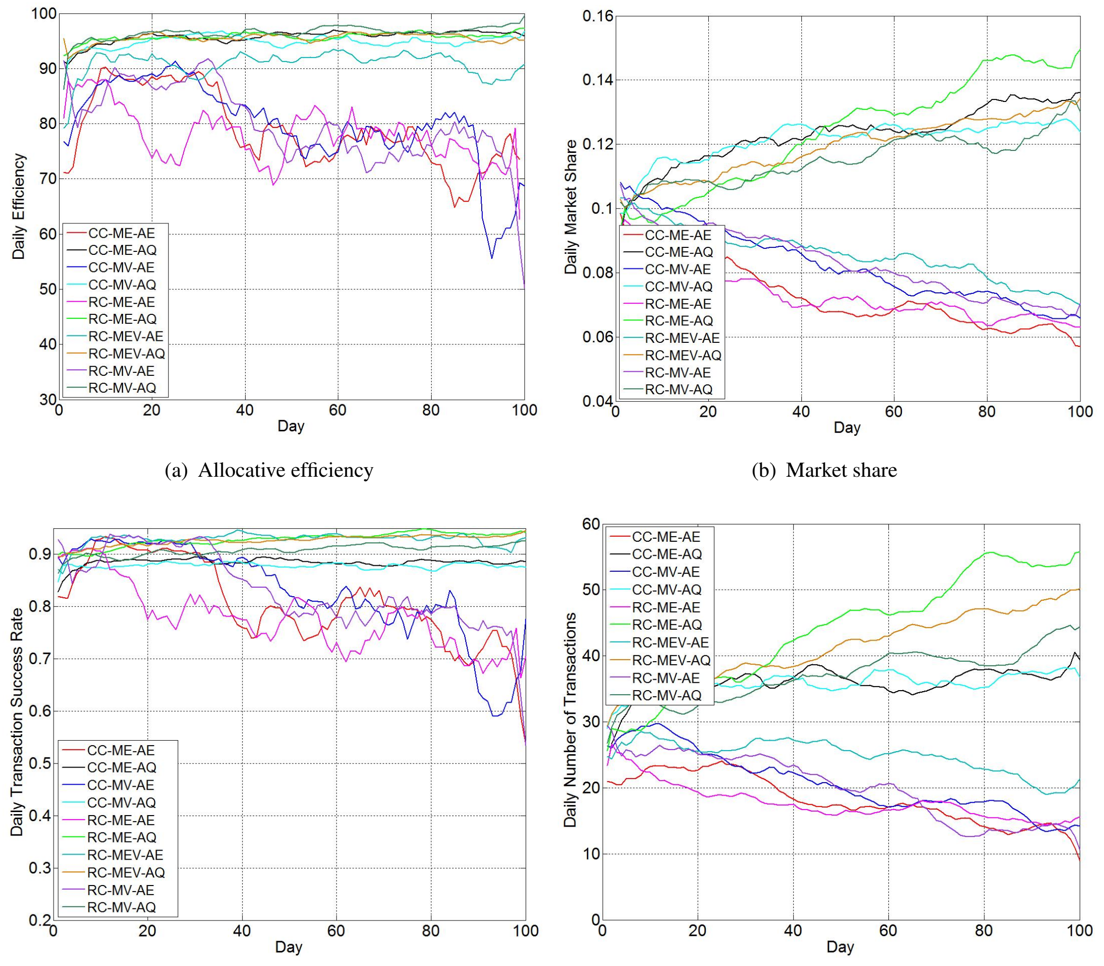

(a) Allocative efficiency (b) Market share

(c) TSR (d) Number of transactions

FIGURE 5.9: Scores of marketplaces in the environment where most traders use ZIP strategy.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	77.695±5.344	0.072±0.006	0.810±0.053	18.216±3.582
CC-ME-AQ	**95.776±0.183**	**0.123±0.004**	0.885±0.008	35.894±2.35
CC-MV-AE	78.430±4.720	0.083±0.006	0.828±0.041	20.836±3.522
CC-MV-AQ	**94.902±0.616**	**0.122±0.005**	0.879±0.009	35.696±3.415
RC-ME-AE	75.682±6.141	0.073±0.005	0.779±0.061	17.802±3.600
RC-ME-AQ	**95.761±0.064**	**0.124±0.005**	**0.929±0.004**	**43.818±2.839**
RC-MEV-AE	90.904±0.830	0.086±0.004	**0.929±0.006**	24.898±2.409
RC-MEV-AQ	**95.512±0.409**	0.119±0.003	**0.927±0.002**	**41.156±1.147**
RC-MV-AE	79.402±4.219	0.083±0.004	0.834±0.045	20.08±2.618
RC-MV-AQ	**96.537±0.239**	0.116±0.003	0.908±0.009	37.088±2.639

TABLE 5.9: Average daily results and corresponding 95% confidence intervals of marketplaces in the environment where most traders use ZIP strategy.

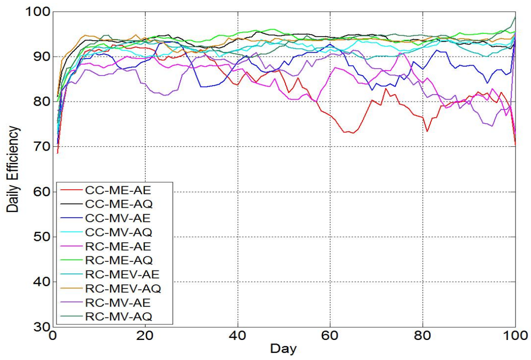

(a) Allocative efficiency

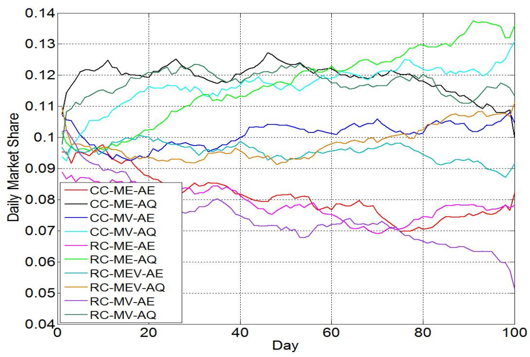

(b) Market share

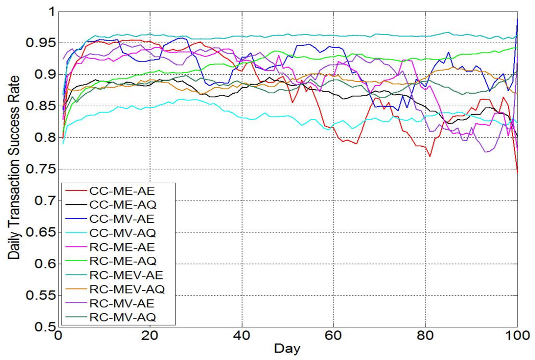

(c) TSR

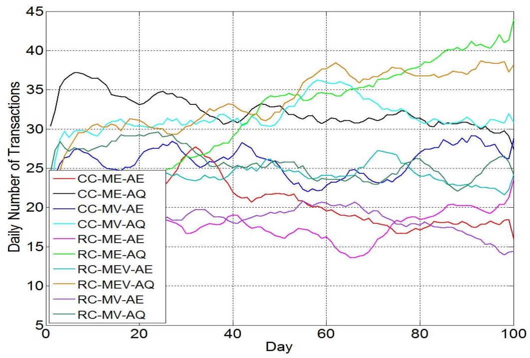

(d) Number of transactions

FIGURE 5.10: Scores of marketplaces in the environment where the numbers of traders using different bidding strategies are equal.

Marketplace	Alloc. Eff. %	Market Share	TSR	Num. of Transactions
CC-ME-AE	83.840±4.131	0.081±0.004	0.881±0.040	20.462±2.773
CC-ME-AQ	**93.498±0.420**	**0.119±0.006**	0.869±0.005	30.294±1.796
CC-MV-AE	87.651±3.332	0.101±0.008	0.914±0.032	29.728±4.086
CC-MV-AQ	91.645±0.551	**0.117±0.004**	0.835±0.014	26.16±2.730
RC-ME-AE	85.079±3.079	0.079±0.004	0.897±0.028	18.206±2.296
RC-ME-AQ	**93.649±0.232**	**0.117±0.002**	0.916±0.005	**31.29±1.250**
RC-MEV-AE	91.311±0.298	0.095±0.003	**0.959±0.003**	**34.130±1.409**
RC-MEV-AQ	**93.331±0.288**	0.098±0.004	0.888±0.007	**35.62±1.178**
RC-MV-AE	85.143±2.352	0.075±0.001	0.896±0.027	16.996±1.175
RC-MV-AQ	**93.047±0.164**	**0.118±0.002**	0.881±0.005	29.248±1.070

TABLE 5.10: Average daily results and corresponding 95% confidence intervals of marketplaces in the environment where the numbers of traders using different bidding strategies are equal.

sactions, and thus can improve TSR. However, since it fails to guarantee traders' profits, some traders may leave the marketplace using the MV matching policy. Therefore, in our competing marketplace, we adopt the MEV matching policy, which is a trade-off between the ME and the MV matching policy. In more detail, in the MEV matching policy, when shouts are far away from the equilibrium price, we use the ME matching policy to maximise traders' profits, and when traders' shouts are within an area close to the equilibrium price, we use the MV matching policy to maximise the number of transactions. By so doing, this policy can guarantee traders' profits and increase the number of transactions, and from Section 5.2 we actually have seen that this policy performs well.

Now we consider the shout accepting policy. This determines the shouts accepted by the marketplace, and thus significantly affects the TSR. If we provide tight restrictions on accepting shouts, our TSR will be improved. However, some intra-marginal traders' shouts will not be accepted by the marketplace, and they will leave the marketplace. Therefore we need to make a trade-off between these. In Section 5.2, we have shown that the quote-beating accepting policy provides a loose restriction on accepting shouts, and the equilibrium accepting policy provides a tight restriction on accepting shouts. Therefore, in our marketplace, when the number of transactions decreases, we will switch to the quote-beating accepting policy. If the number of transactions is at a good level, we will use the equilibrium accepting policy to improve TSR.

Finally, for the pricing policy,, market theory has indicted that when the transactions happen at the equilibrium price, the optimal allocative efficiency is reached (see Section 2.2). Therefore, we use the equilibrium price as the transaction price in the CAT competition. In more detail, if the equilibrium price is higher than the matched ask and is lower than the matched bid, we set the equilibrium price as the transaction price; otherwise we set the transaction price to the bid or ask which is closest to the equilibrium price. Note that in this thesis, we adopt k-pricing policy in the analysis, which is different from what we actually used in the CAT competition. The reason is that k-pricing policy can be easily represented in a mathematic way and it also has been shown to be highly efficient in the literature (Phelps et al., 2003). Furthermore, in a marketplace with symmetric demand and supply curve, we can expect that the transaction price set by k-pricing policy with $k = 0.5$ is approximately equal to the equilibrium price.

Now we have introduced our design of the market policies for the CAT competition. In Section 5.5, we will describe how these policies performed in 2010 CAT competition. First, however, we will detail the design of the charging policy.

5.4 Charging Policy Design

After detailing the market policies for the CAT competition, in this section we discuss the design of our charging policy, which is used to determine fees charged to traders. As we discussed before (see Section 1.2), there exists a conflict between making profits and attracting (or keeping) traders. In particular, the marketplace charging higher fees may obtain high market profits in the short term, but will lose traders in the long term, and so will receive less profits. Thus, a good charging policy should be able to make high market profits, while still maintaining the number of traders at a good level. In particular, in this section, we will design an *adaptive* charging policy that charges market fees based on market conditions.

In so doing, we would like to use the insights from our analysis in Chapters 3 and 4. However, we cannot directly use these results because the setting is different in a number of respects[7]. Firstly, for traders' bidding behaviour, as we discussed in Section 5.1.2, traders using the equilibrium bidding strategies shade their shouts less than when using GD or ZIP strategy, and more than when using ZI-C or RE strategy. Furthermore, in Chapters 3 and 4, although both EGT and FP are repeated learning approaches, the game we analysed is a one-shot game (i.e. restricted in one trading round). In the CAT competition, there are several trading days which includes several trading rounds, and thus it is a repeated game. Finally, in the CAT competition, different marketplaces have different charging policies, and we cannot guarantee that other marketplaces also use the equilibrium charging strategies. Given this situation, we cannot directly use the equilibrium charging strategies we have derived. However, as we will show, a number of the general insights from the equilibrium analysis are still useable. In the following, we first determine what types of fees are effective at making profits and keeping traders. Then we decide at what point in the game to charge fees and how much to charge.

In the CAT competition, competing marketplaces can charge five different types of fees: registration, information, shout, transaction and profit fees (see Section 2.5.1). We now discuss each of these in turn, and determine what types of fees are the most effective. A registration fee is charged to all traders that register with a marketplace. When such a fee is charged, because extra-marginal traders usually cannot trade in the marketplace, their profits will be negative. Thus these traders will leave the marketplace. As extra-marginal traders leave the marketplace, the market share will decrease. Moreover, when a high registration fee is charged, some intra-marginal traders whose limit/cost prices are close to the equilibrium price may also leave the marketplace because of negative profits. This is consistent with our analysis in Section 4.4, where we found that charging registration fees cause traders to leave the marketplace quickly. The information fee is charged only to GD and ZIP traders since only these traders need information provided by the marketplaces to generate shouts. Therefore, this is equivalent to a registration fee, but one that only applies to GD and ZIP traders. A shout fee will be charged when traders successfully place shouts. Like the registration fee, this may drive extra-marginal traders to leave the marketplace when shouts placed by extra-marginal traders are accepted by the marketplace, and thus decreases the market share. We can see that the above three types of fees are charged to traders no matter whether they have made any profits (i.e. ex ante fees, see Section 1.1). This causes some traders to have negative profits and leave the marketplace. Therefore, all these three types of fees are less effective in either making profits or keeping traders. In contrast, transaction and profit fees are charged only when traders make transactions, and thus make profits (i.e. ex post fees, see Section 1.1). Therefore, in such cases, traders'

[7]We have to operate the theoretical analysis on the restricted setting to provide basic insights. In the future, we will aim to close the gap between the theoretical and practical aspects.

profits are usually positive. Furthermore, charging these two types of fees will not drive extra-marginal traders to leave the marketplace, which will benefit the market share. This is consistent with our analysis in Section 4.4, where we found that the marketplace charging transaction fees can maintain traders longer than the marketplace charging registration fees, and the marketplace charging profit fees can keep traders even though it charges very high profit fees. Furthermore, from Section 4.4, we found that charging transaction fees is effective in making profits and charging profit fees cannot guarantee high market profit when traders can shade their shouts a lot. However, we should note that when traders cannot significantly shade their shouts (e.g. ZI-C or RE traders), i.e. they cannot effectively hide their actual profits, charging profit fees will guarantee high market profits. Now, given this analysis, in terms of making profits and keeping traders, we decide to charge two types of fees: transaction and profit fees. Note that in Section 4.4, we found that the transaction price percentage fee is the most effective in making profits and keeping traders. However, this type of fees is not allowed in the CAT competition.

After determining what types of fees to charge to traders, we now describe when and how much to charge. From Section 3.2.3.3, we obtained the insight that, initially, the marketplace should charge low (or even zero) fees to attract traders, and when it obtains a larger market share, it can charge higher fees, but still keep traders. In our charging policy, we adopt this insight. However, we should note that in Section 3.2.3.3, we also found that, when there is a large number of traders, or traders are able to explore other marketplaces to search for cheaper alternatives, it is difficult for the marketplace to keep traders even though it already has a larger market share. In the context of CAT competition, there are a large number of traders (400 traders in recent years' competition), and traders with an ϵ-greedy exploration strategy are able to search for the cheaper marketplace. Thus it is difficult for the marketplace to keep traders when it charges higher fees in the CAT competition. We address this problem by reducing the fees whenever the transaction share falls below a certain threshold, thereby repeating the above problem. Here the transaction share is the number of transactions made in this marketplace as a percentage of the total number of transactions made in all marketplaces. The reason that instead of directly looking at market share, we look at the transaction share is as follows. As we discussed above, our marketplace charges fees after traders make transactions. However, a large number of traders does not always mean a large number of transactions since it may happen that most traders are extra-marginal traders. Thus it is better for determining whether charging fees or not based on transaction share, instead of market share.

Now we discuss how to calculate the marketplace's transaction share in the CAT competition. In order to do this, we need to know the number of transactions in other marketplaces. In the JCAT platform, the marketplace can subscribe to other marketplaces to obtain this information, but an information fee needs to be paid to subscribed marketplaces. This payment (incurred by an information fee) will result in a "lost profit share" for the marketplace, which is the proportion of total information fees to the total profits of all marketplaces. Thus, in the beginning of each competition day, we need to decide which marketplace we need to subscribe to. In more detail, we do not require information about marketplaces which have historically low transaction share since low transaction share indicates that these marketplaces perform badly. Furthermore, when the "lost profit share" is less than a predetermined threshold, then the marketplace purchases information from all marketplaces. However, if the "lost profit share" is higher than this threshold, then the specialist does not purchase information from marketplaces with high information fees and low historical transaction shares, so as to keep the "lost profit share" below the threshold. However, these marketplaces may perform better later, and therefore it is necessary to relax the threshold regularly in order to obtain information from these marketplaces. After obtaining information about transactions from other marketplaces, we can calculate the marketplace's transaction share directly.

Then our charging policy proceeds as follows. In the first few days, our marketplace charges no fees in order to attract traders, and thus build up its transaction share. When the marketplace obtains a large transaction share (i.e. reaches a predetermined threshold), it then starts to charge transaction and profit fees to extract profits from traders. With increasing fees, the marketplace becomes less attractive to potential traders (since traders are able to search for the cheaper marketplace), and then gradually, the effect of increased fees decreases the number of traders, and thus the transaction share. At this point, the marketplace decrease its fees and goes back to building up its transaction share back to a predetermined threshold. Once it reaches this target, it has a sufficient number of transactions, and can charge high fees to make profits again. By following this process, the marketplace will *adapt* fees according to its transaction share, and can make high profits while maintain the transaction share at a reasonably good level.

So far we have established when to charge fees. Now, in order to set the level of fees, we use the concept of target profit. This is the profit that our marketplace attempts to extract from traders on the current day. Intuitively, we can see that the target profit depends on the number of transactions made in the marketplace. If more transactions are made, the marketplace can extract more profits from traders. In more detail, we consider that the target profit share is approximately equal to the transaction share of the marketplace, and then the target profit will be equal to the product of the marketplace's transaction share and the estimated historical total profits extracted from traders by all marketplaces. After calculating the target profit, we now describe how to set transaction and profit fees to obtain such a profit.

Since we charge two different types of fees, we need to split the target profit to two parts, each of which is extracted from charging transaction fee and profit fee respectively. We now discuss how to split the target profit. To this end, the shaded area in Figure 5.11 shows the observed profits of traders. Here, the observed profits are the sum of the differences between the matched bids and asks. Note that the observed profits are not equal to the actual profits of traders unless traders adopt a truth-telling strategy. If traders use the ZI-C or RE strategies, the observed profits are relatively close to the actual profits, compared to that of traders using ZIP or GD strategies. Given the target profit, when only a profit fee is charged, the target profit is equal to the size of the dark triangular area. When only a transaction fee is charged, the target profit is equal to the size of the rectangular area surrounded by the red line. From the figure, we can see that, traders whose shouts are far away from the equilibrium price will pay more when a profit fee is charged

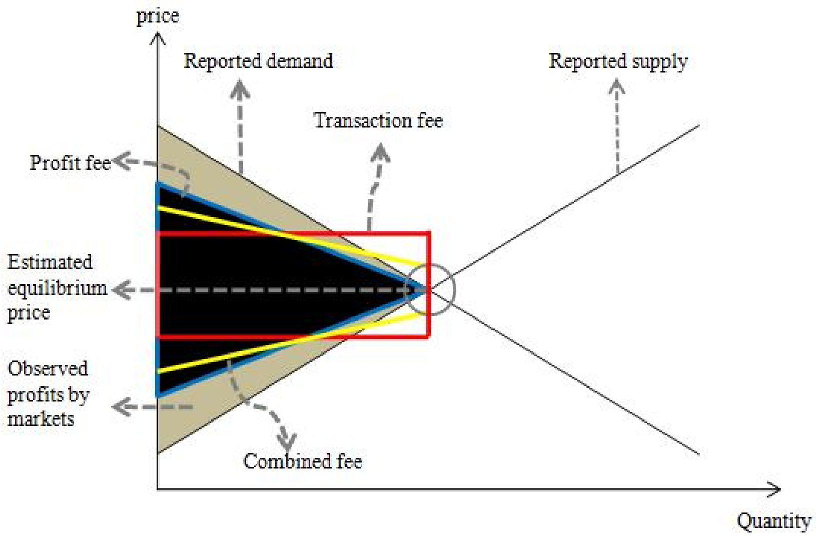

FIGURE 5.11: Trade-off between charging a profit fee and a transaction fee. Red, blue and yellow lines represent different types of fees charged to traders respectively: transaction fee, profit fee, and a combination of transaction and profit fee.

than when a transaction fee is charged; and traders whose shouts are close to the equilibrium price pay more when a transaction fee is charged than that when a profit fee is charged. Thus in the design of a charging policy, we need to make a trade-off between these two fees, as showed by the yellow line in Figure 5.11, which we call the combined fee. In more detail, initially, we assume that more than half of the target profit is obtained from charging a transaction fee since charging a transaction fee can guarantee market profit in any cases (even though traders shade shouts a lot), and the rest from a profit fee. As the game proceeds, we adjust the balance as follows. If most shouts are close to the equilibrium price, i.e. most traders shade their shouts a lot, charging a profit fee cannot guarantee a target profit, and thus we extract more from charging transaction fees. On the other hand, when less shouts are close to the equilibrium price, charging a profit fee can guarantee the target profit. We then extract relatively more from charging profit fees. However, in this case, traders with shouts far away from the equilibrium price will pay more when profit fees are charged, and thus may leave the marketplace. To avoid this, we assume that at least half target profit is extracted from charging transaction fees.

5.4.1 Evaluation of the Charging Policy

Now that we have presented our adaptive charging policy, in this section, we evaluate it against a number of the charging policies used in recent years' CAT competition. Note that all competition entrants only published their binary codes of marketplace implementation, and thus we do not know exactly what charging policies other marketplaces used in the competition. We can only observe the fees charged to traders from the competition log files. However, by investigating the log files from recent years, we found that the marketplaces that perform well only charge profit fees and these fees are fairly constant. Given these, we consider three marketplaces, which charge fixed 5%, 10% and 15% profit fees respectively (which is similar to what most of the marketplaces did in the competition). Although very few marketplaces charge transaction fees in the CAT competition, according to our analysis in Section 4.4, this type of fees is effective in making profits and keeping traders. Therefore, we consider three marketplaces charging fixed 0.5, 1 and 1.5 transaction fees respectively, which correspond more or less to the above profit fees in terms of the absolute payments incurred by them[8]. Furthermore, we consider a marketplace charging a fixed 0.1 registration fee as a representative marketplace charging other types of fees. For all these marketplaces, we assume that they begin to charge fees on day 21 (i.e. in the first 20 days, they charge no fees to build up their market share). In terms of the market policies, we consider that all marketplaces adopting the default market policies provided by JCAT platform, which are RC-ME-AQ and k-pricing policy with $k = 0.5$ (in Section 5.2, we have shown that these market policies perform well).

We now run simulations to evaluate our charging policy against the above fixed policies. The experimental setup is as follows. Each experiment runs for 200 days with 10 rounds per day and 1 second per round. There are 200 buyers and 200 sellers. The private values of all traders are independently drawn from a uniform distribution between 50 and 150, and each trader is allowed to buy or sell up to 3 goods per day. We will run our simulations in different environments with different bidding strategies since the effectiveness of a charging policy is affected by traders' strategies. For each environment, the experiment is repeated 40 times. Firstly, we consider the environment where all traders use ZI-C bidding strategy. We evaluate the performance in terms of market share, profit share, the sum of weighted market share and profit share with 0.5 weight on each and the number of transactions. The results are shown in Figure 5.12 and Table 5.11, from which, we can see that our charging policy can significantly outperform other marketplaces. From Figures 5.12(a) and 5.12(b), we can see that once competing marketplaces begin to charge fees, our marketplace will attract traders and make profits at a good level. In the whole competition, we can see that our market share is maintained at a good level, and our profit share is quite high. This is because our charging policy determines the target profit according to the changes of the transaction share, and thus will not set unreasonable fees to extract profits. Moreover, by extracting the target profit from charging both transaction and profit fees, our marketplace is better on keeping traders than those marketplaces only charging a profit fee or a transaction fee. Thus our marketplace can obtain a good profit share while maintaining traders at a good level. Furthermore, we also find that the market share of the marketplace charging a registration fee is very low since no extra-marginal traders stay in this marketplace. We then run experiment in the environment of all traders using RE strategy, and the results are shown in Figure 5.13 and Table 5.12. We also find that the marketplace using our charging policy obtains the highest market share and profit share, and it outperforms other marketplaces significantly.

We further run experiments in the environments where all traders use GD and ZIP strategies

[8] When traders' private values are drawn from a uniform distribution between 50 and 150, we roughly consider that the average of the buyers' bids is around 110 and the average of the sellers' asks is around 90. Given that the transaction prices are set as 100, the absolute payments incurred by 5%, 10% and 15% profit fees are 0.5, 1 and 1.5 respectively.

respectively. The results are shown in Figure 5.14, Table 5.13 and Figure 5.15, Table 5.14 respectively. We can see that our marketplace outperforms other marketplaces. Furthermore, since traders shade their shouts a lot (see Figure 5.1(d) and 5.1(c)), we find that the marketplaces charging profit fees can maintain traders at a higher level compared to marketplaces charging transaction fees (see Figures 5.14(a), 5.15(a) and Tables 5.13 and 5.14), and the marketplaces charging transaction fees can obtain profit share at a high level compared to marketplaces charging profit fees (see Figures 5.14(b), 5.15(b) and Tables 5.13 and 5.14).

Finally, we evaluate our charging policy in an environment where traders use different bidding strategies. Specifically, we assume that there are 100 ZI-C traders, 100 RE traders, 100 GD traders and 100 ZIP traders. Experimental results are shown in Figure 5.16 and Table 5.15. Again, we can see that the marketplace adopting our charing policy has the highest market share and profit share. This shows that our charging policy is also effective in making profits and keeping traders when traders adopt heterogeneous bidding strategies.

5.5 The 2010 CAT Competition

Now we have introduced our design of market policies and the charging policy. In 2010 CAT competition, on the first day, we entered our marketplace using policies not based on the results of this thesis (as the software was not available). On the second and third days of the competition, which took place a while later, we did use market policies and the charging policy designed in this thesis. The competition results showed that our marketplace performed well in the last two days' competition, obtaining the first position on the second day and the second position on the third day among nine entrants. In the following, we describe how our marketplace performed on the second and third day.

The competition results of the second day and the third day are shown in Figure 5.17 and Table 5.16 and Figure 5.18 and Table 5.17 respectively. On the second day, we obtained the first position, and on the third day, we obtained the second position. In more detail, we can see that our competing marketplace obtained good scores on the sum of daily market share, the sum of daily profit share and the sum of daily TSR. The high market share shows that the timing, matching and pricing policies used by our competing marketplace are highly effective in terms of guaranteeing traders' profits, and thus can attract and keep traders. The high TSR shows that our shout accepting policy is highly effective to reject shouts which are unlikely to be matched, and the high market share also shows that the shout accepting policy does not reject shouts that can make transactions. Furthermore, from the high scores of profit share and market share, we can see that our marketplace using the adaptive charging policy designed in this chapter makes a high market profit while maintaining the number of traders at a good level. This shows that our design of the charging policy, which adapts the fees according to the transaction share and makes a trade-off between only charging a profit fee and only charging a transaction fee, is highly effective. However, although our competing marketplace performed well in the competition, it

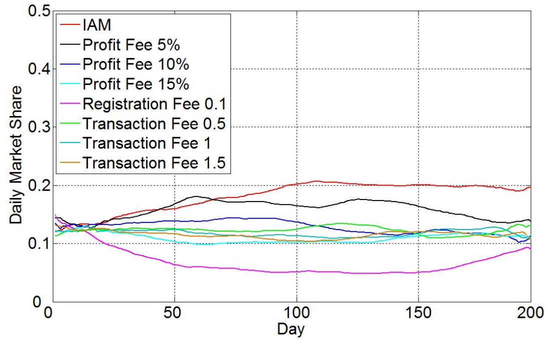

(a) Market share

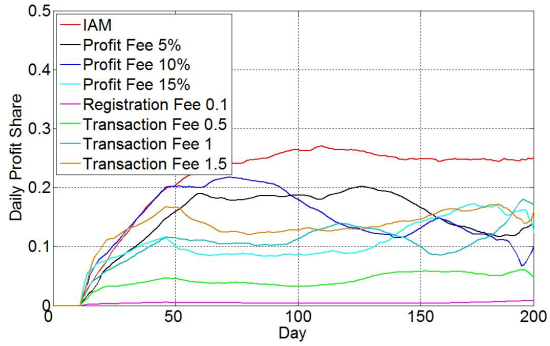

(b) Profit share

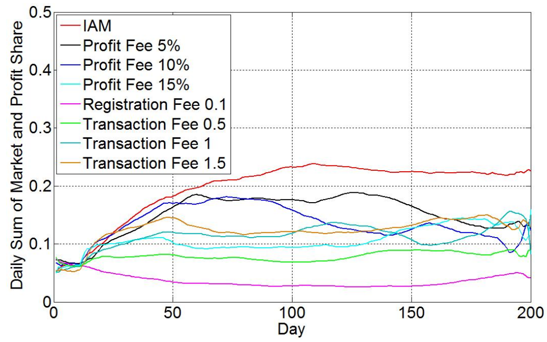

(c) Sum of weighted market share and profit share

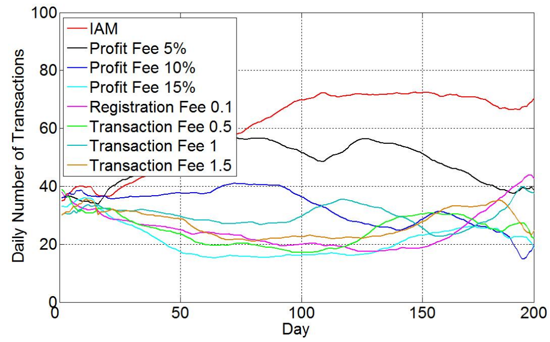

(d) Number of transactions

FIGURE 5.12: Evaluation of charging policy with ZI-C strategy.

Marketplace	Market Share	Profit Share	Sum of Weighted Market and Profit Share	Num. of Transactions
IAM	0.181±0.002	0.223±0.002	0.202±0.002	54.44±1.248
PF 5%	0.158±0.003	0.13±0.004	0.144±0.003	48.966±2.382
PF 10%	0.128±0.002	0.143±0.002	0.136±0.001	32.644±1.243
PF 15%	0.11±0.002	0.107±0.003	0.108±0.002	21.419±1.682
RF 0.1	0.067±0.002	0.004±0.001	0.036±0.002	25.017±1.535
TF 0.5	0.124±0.001	0.042±0.002	0.083±0.002	29.895±1.212
TF 1	0.117±0.002	0.104±0.003	0.111±0.002	27.989±1.741
TF 1.5	0.115±0.002	0.141±0.002	0.128±0.002	24.015±1.257

TABLE 5.11: Evaluation of charging policy with ZI-C strategy. PF: Profit Fee. RF: Registration Fee. TF: Transaction Fee.

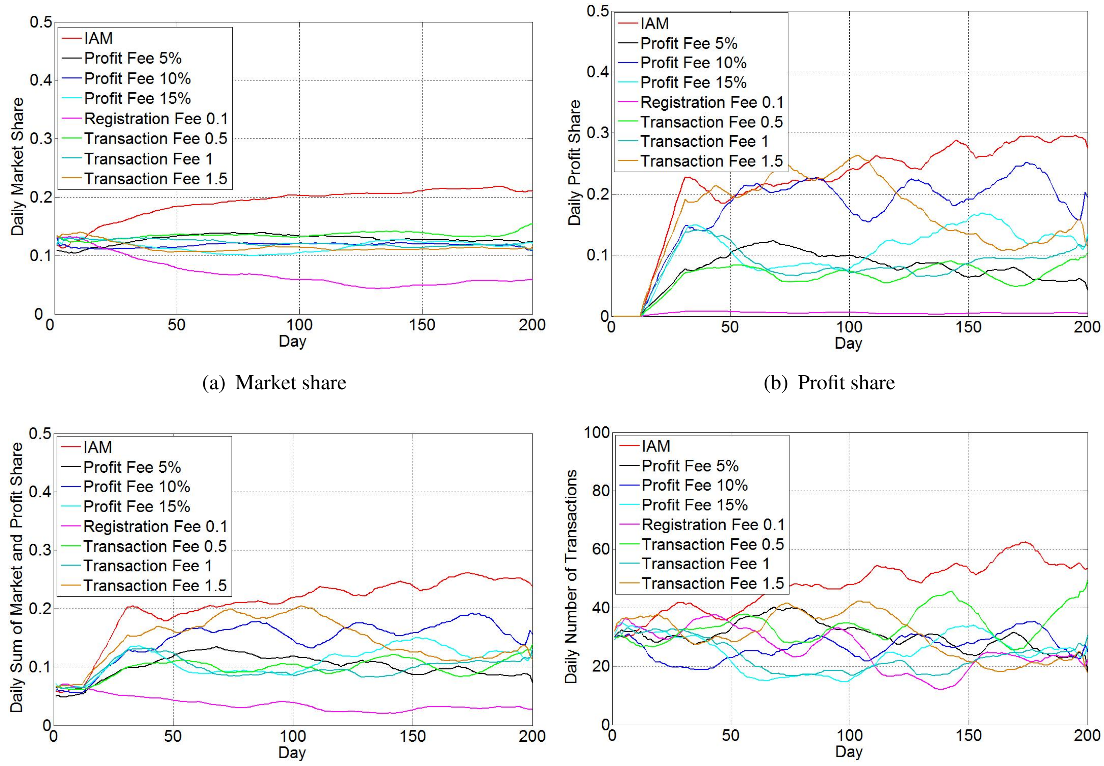

(a) Market share

(b) Profit share

(c) Sum of weighted market share and profit share

(d) Number of transactions

FIGURE 5.13: Evaluation of charging policy with RE strategy.

Marketplace	Market Share	Profit Share	Sum of Weighted Market and Profit Share	Num. of Transactions
IAM	0.193±0.002	0.211±0.002	0.202±0.002	46.586±1.559
PF 5%	0.125±0.002	0.077±0.002	0.101±0.002	32.847±1.672
PF 10%	0.121±0.002	0.185±0.004	0.153±0.004	26.155±1.91
PF 15%	0.116±0.003	0.106±0.004	0.111±0.003	22.817±2.223
RF 0.1	0.069±0.002	0.005±0.002	0.037±0.002	25.93±1.551
TF 0.5	0.133±0.003	0.064±0.002	0.098±0.003	33.761±2.56
TF 1	0.124±0.002	0.083±0.002	0.104±0.002	30.373±1.49
TF 1.5	0.114±0.002	0.164±0.004	0.139±0.004	23.761±1.653

TABLE 5.12: Evaluation of charging policy with RE strategy.

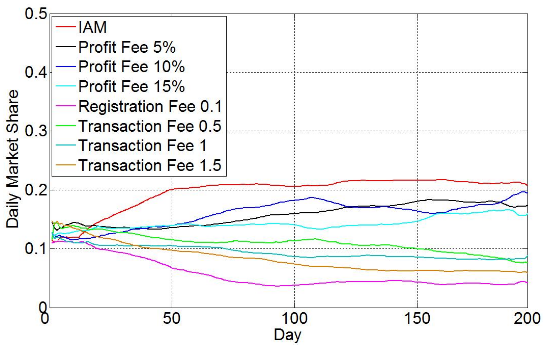

(a) Market share

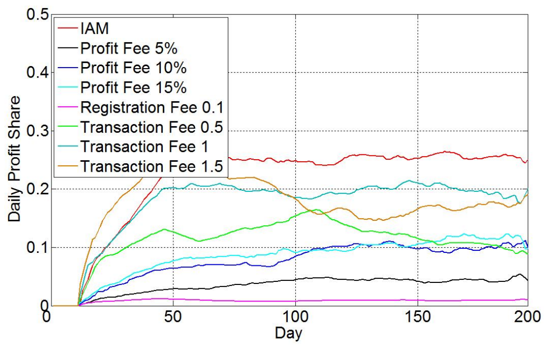

(b) Profit share

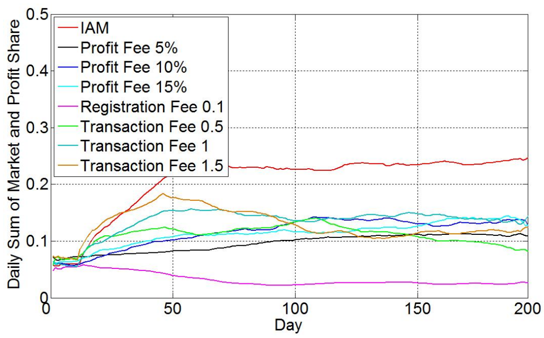

(c) Sum of weighted market share and profit share

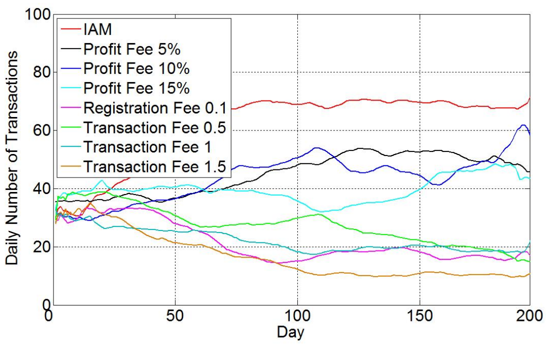

(d) Number of transactions

FIGURE 5.14: Evaluation of charging policy with GD strategy.

Marketplace	Market Share	Profit Share	Sum of Weighted Market and Profit Share	Num. of Transactions
IAM	0.198±0.002	0.215±0.002	0.206±0.002	55.693±1.965
PF 5%	0.159±0.004	0.04±0.003	0.099±0.003	48.231±2.178
PF 10%	0.152±0.004	0.074±0.002	0.113±0.004	44.345±2.249
PF 15%	0.144±0.004	0.085±0.002	0.114±0.004	39.962±1.992
RF 0.1	0.056±0.002	0.008±0.001	0.032±0.002	21.477±1.482
TF 0.5	0.109±0.002	0.113±0.002	0.111±0.002	27.375±1.619
TF 1	0.093±0.002	0.181±0.002	0.137±0.002	22.017±1.51
TF 1.5	0.082±0.002	0.178±0.002	0.13±0.002	16.851±1.337

TABLE 5.13: Evaluation of charging policy with GD strategy.

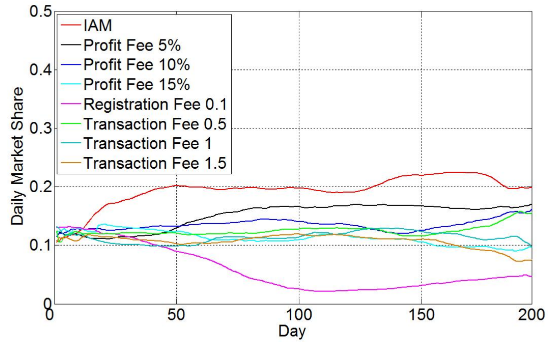

(a) Market share

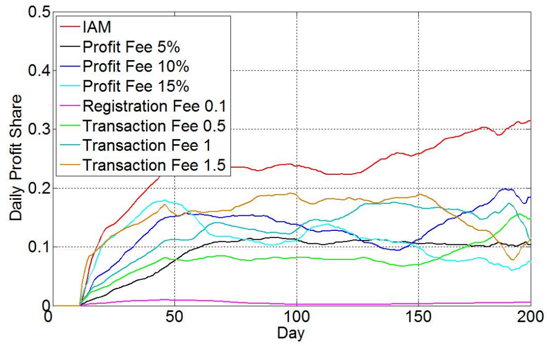

(b) Profit share

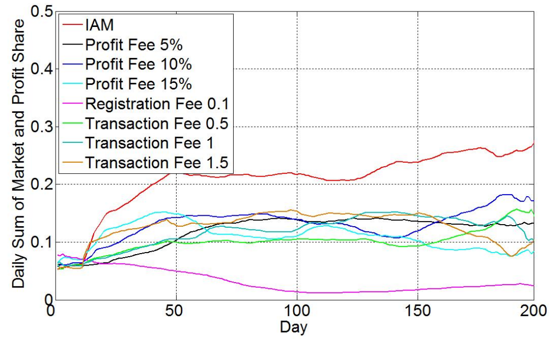

(c) Sum of weighted market share and profit share

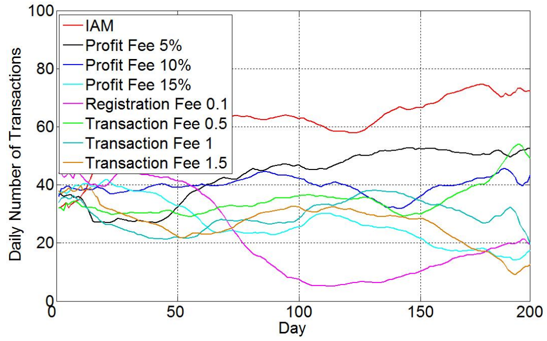

(d) Number of transactions

FIGURE 5.15: Evaluation of charging policy with ZIP strategy.

Marketplace	Market Share	Profit Share	Sum of Weighted Market and Profit Share	Num. of Transactions
IAM	0.196±0.002	0.201±0.002	0.198±0.002	54.234±1.949
PF 5%	0.149±0.004	0.078±0.002	0.113±0.004	48.084±2.233
PF 10%	0.137±0.004	0.123±0.004	0.130±0.004	41.454±2.033
PF 15%	0.113±0.002	0.107±0.004	0.11±0.002	26.876±1.594
RF 0.1	0.059±0.002	0.005±0.001	0.032±0.002	22.295±1.472
TF 0.5	0.125±0.004	0.087±0.002	0.106±0.003	34.72±1.864
TF 1	0.113±0.002	0.134±0.004	0.123±0.004	29.294±1.702
TF 1.5	0.107±0.002	0.16±0.004	0.134±0.004	26.258±1.596

TABLE 5.14: Evaluation of charging policy with ZIP strategy.

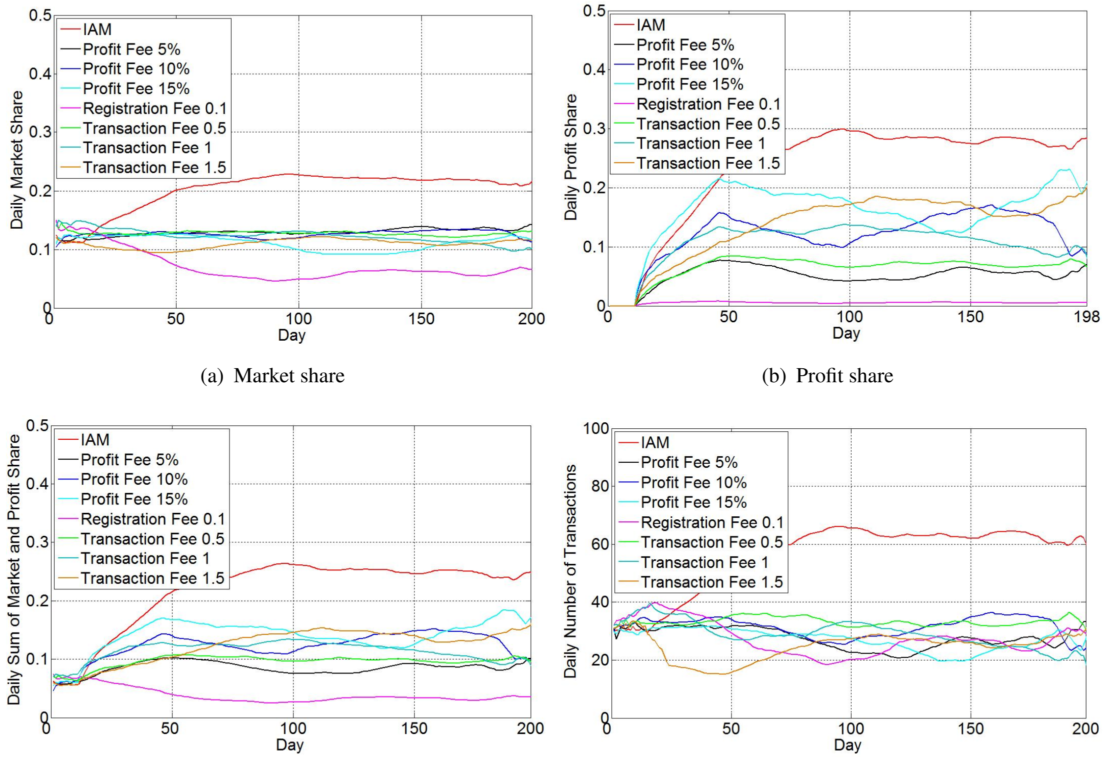

(a) Market share

(b) Profit share

(c) Sum of weighted market share and profit share

(d) Number of transactions

FIGURE 5.16: Evaluation of charging policy with heterogeneous strategies.

Marketplace	Market Share	Profit Share	Sum of Weighted Market and Profit Share	Num. of Transactions
IAM	0.203±0.002	0.222±0.002	0.212±0.002	53.079±1.702
PF 5%	0.126±0.004	0.052±0.002	0.089±0.002	37.523±1.666
PF 10%	0.122±0.004	0.124±0.004	0.123±0.004	31.707±1.776
PF 15%	0.119±0.004	0.178±0.006	0.148±0.004	25.091±2.251
RF 0.1	0.071±0.002	0.005±0.002	0.038±0.002	27.49±1.637
TF 0.5	0.127±0.004	0.066±0.002	0.096±0.002	33.341±2.117
TF 1	0.122±0.002	0.108±0.004	0.115±0.004	29.382±1.668
TF 1.5	0.109±0.002	0.138±0.004	0.124±0.004	23.707±1.588

TABLE 5.15: Evaluation of charging policy with heterogeneous strategies.

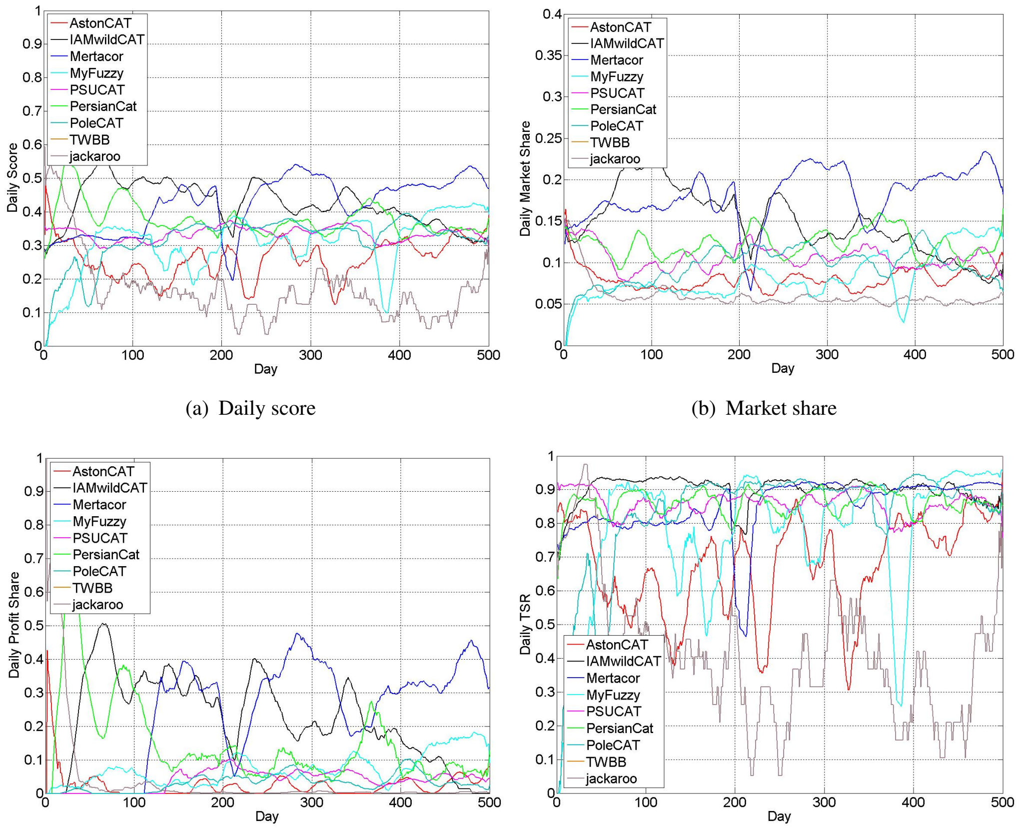

(a) Daily score (b) Market share

(c) Profit share (d) TSR

FIGURE 5.17: Competition on the second day.

Rank	Specialist	Total Score	Total Market Share	Total Profit Share	Total TSR
1	IAMwildCAT	198.228	69.923	110.582	414.199
2	jackaroo	193.008	56.508	83.266	439.035
3	Mertacor	191.679	82.001	106.06	386.393
4	PoleCAT	174.621	56.073	72.319	395.481
5	PSUCAT	153.013	48.203	18.368	392.597
6	TWBB	148.324	44.928	14.953	384.151
7	MyFuzzy	139.166	36.469	29.93	350.114
8	AstonCAT	115.685	35.906	8.229	303.241
9	PersianCat	73.765	27.026	13.31	182.512

TABLE 5.16: Competition result on the second day in 2010 CAT competition

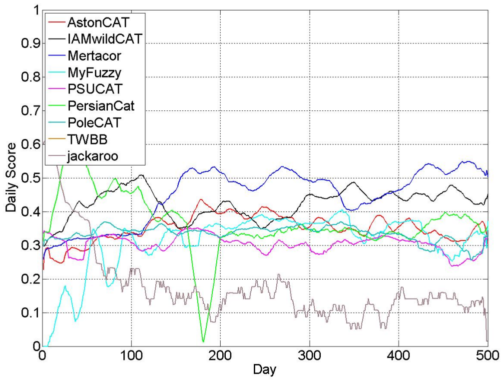

(a) Daily score

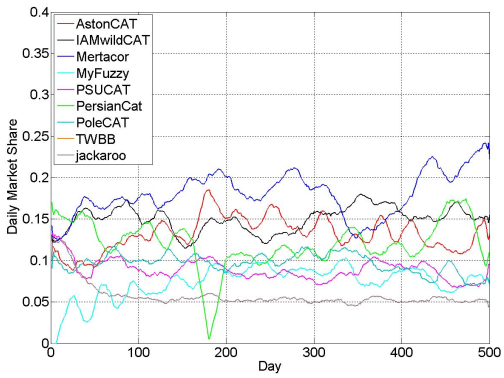

(b) Market share

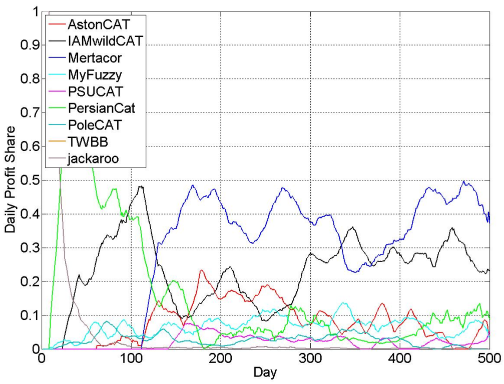

(c) Profit share

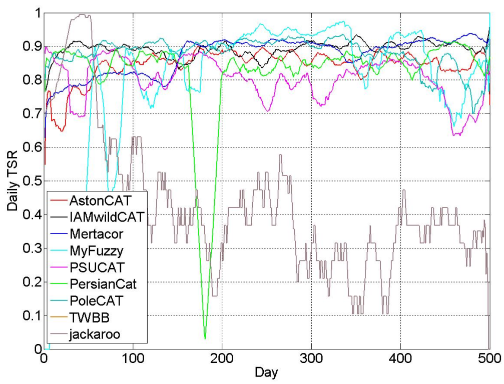

(d) TSR

FIGURE 5.18: Competition on the third day.

Rank	Specialist	Total Score	Total Market Share	Total Profit Share	Total TSR
1	Mertacor	208.576	83.289	135.125	406.367
2	IAMwildCAT	198.421	69.988	109.403	415.548
3	jackaroo	172.217	46.303	39.357	431.686
4	PoleCAT	170.951	55.662	75.222	381.917
5	AstonCAT	164.445	61.543	38.651	393.029
6	TWBB	157.257	45.567	13.204	431.686
7	MyFuzzy	148.131	35.082	31.325	377
8	PSUCAT	141.819	41.291	9.319	375.029
9	PersianCat	79.7	26.26	13.368	200.297

TABLE 5.17: Competition result on the third day in 2010 CAT competition

failed to beat Mertacor (which is the winner of this year's competition) on the third day. This means that our design of market policies and the charging policy can be improved.

5.6 Summary

In this chapter, we first described how traders select marketplaces and submit shouts in the CAT competition. We then ran experiments to analyse how different market policies affect marketplaces' performance in different environments with different bidding strategies. This analysis addresses our research challenge 3 of analysing market policies (see Section 1.2). From this analysis, we obtained several insights, and we further used these insights to design market policies for the CAT competition. For example, in the first round of each trading day, we adopt the continuous clearing in order to provide information for traders to improve their shouts, and we switch between quote-beating accepting and equilibrium accepting policies in order to keep traders and improve TSR. After designing market policies, we then used insights from Chapters 3 and 4 to design a novel charging policy, which will charge transaction and profit fees, and adapt fees according to transaction share. We further evaluated it in the context of CAT competition. Finally, we showed that our design of market policies and charging policy performs well in the actual CAT competition (ranked the first on the second day and the second on the third day). This work then addresses our research challenge 5 of designing an effective competing double auction marketplace (see Section 1.2).

Chapter 6

Conclusions and Future Work

The double auction, a highly efficient market mechanism, has been widely used by both traditional and online exchanges. In today's economy, these marketplaces increasingly need to compete with each other to attract traders and make profits by appropriately charging fees to participating traders. Therefore, it is necessary to design effective market policies and charging strategies that can operate effectively in this situation. To this end, in this thesis, we analysed how double auction marketplaces can be designed to compete with each other in an effective way and then use the insights from this analysis to design an effective competing marketplace agent. In so doing, we have made several contributions (summarised below in Section 6.1) to the state of the art. Thereafter, we outline the directions for future work in this area in Section 6.2.

6.1 Research Summary

The design of a competing double auction marketplace primarily consists of setting the market policies and a charging strategy. In this thesis, we mainly focused on how competing marketplaces should set their fees (i.e. the charging strategy), since this is a significant determinant that affects traders' market selection and the marketplaces' profits. Now, because such charging strategies are affected by both the traders' market selection and bidding strategies, we first analysed the traders' strategies and then the marketplaces' charging strategies. Since the optimal behaviour of a trader in terms of selecting marketplaces and submitting shouts depends on the behaviour of other traders and marketplaces, and the optimal behaviour of a marketplace in terms of setting fees depends on the behaviour of traders and other marketplaces, we used game theory to analyse the equilibrium strategies for traders and marketplaces. In addition to analysing the charging strategies, we further analysed how different market policies affect the performance of competing marketplaces in the CAT competition. Finally, based on the insights from analysing the charging strategies and market policies, we designed an effective competing double auction marketplace, and entered it into 2010 CAT competition, where it was very successful.

In more detail, firstly, in Chapter 3, we used game theory to analyse equilibrium market selection strategies for traders and equilibrium charging strategies for marketplaces in the setting with discrete trader types. We also assumed that traders use a truth-telling bidding strategy and traders can only enter one marketplace at a time (i.e. single-home trading). This work addressed research challenges 1 (analysing market selection strategies) and 4 (analysing charging strategies), and it is the *first* theoretical work on analysing equilibrium strategies for marketplaces and traders in the context of multiple competing double auction marketplaces. Specifically, we first analysed the equilibrium market selection strategies for traders for a given fee system. We used game theory to derive the equilibrium market selection strategies analytically and used evolutionary game theory to investigate the dynamics of traders' strategies. In so doing, we showed which equilibrium traders are more likely to converge to. Furthermore, we found that, when the same type of fees are charged by two marketplaces, all the traders will congregate in one marketplace. However, when different types of fees are allowed (registration fees and profit fees), competing marketplaces are more likely to co-exist in equilibrium. Furthermore, we found an interesting phenomenon that sometimes all the traders eventually migrate to the marketplace that charges higher fees, when this marketplace initially has a larger proportion of the traders. We further analysed this phenomenon in detail. This analysis provided us the insight into the charging strategy, which is that, firstly, a marketplace should lower its fees to attract or maintain traders, and after obtaining an advantageous position, the marketplace can then increase its fees while still keeping traders. Based on the analysis of traders' equilibrium market selection strategies, we analysed the equilibrium charging strategies of the marketplaces using two different approaches. In the first, we derived the equilibrium charging strategies by a static analysis. However, this approach did not consider the interaction between the traders' and the marketplaces' strategies. We then tackled this limitation by using a co-evolutionary approach to analyse how competing marketplaces dynamically set fees, while taking into account the dynamics of the traders' market selection strategies. In so doing, we found that two initially identical marketplaces undercut each other, and they will eventually charge the minimal fee that guarantees positive market profits for them. Furthermore, we also extended the co-evolutionary analysis of the marketplaces' charging strategies to more general cases. Specifically, we analysed how an initially disadvantaged marketplace with an adaptive charging strategy can beat an initially advantaged one with a fixed charging strategy, and how competing marketplaces evolve their charging strategies when different types of fees are allowed.

The work in Chapter 3 was restricted to the setting with discrete trader types and assumed that traders adopt a simple, truth-telling bidding strategy. In Chapter 4, we addressed these shortcomings by considering continuous trader types and analysing both the equilibrium market selection and bidding strategies for traders. Moreover, we considered two more types of fees: transaction and transaction price percentage fees. Furthermore, we extended this analysis to settings with different trading environments and with different good properties. By considering these additional factors, we used a fictitious play algorithm to analyse how traders select marketplaces and submit shouts, and how competing marketplaces set fees in equilibrium. This work addressed research challenges 1 (analysing market selection strategies), 2 (analysing bidding strategies) and 4 (analysing charging strategies). In more detail, we first analysed traders' equilibrium bidding strategies in a single marketplace and analysed the effect of different types of market fees on the traders' equilibrium bidding strategies. We observed that registration fees cause a bigger range of traders not to choose the marketplace; profit fees cause traders to shade their shouts more; and transaction price percentage fees cause sellers to shade relatively less than buyers' shading. This is the *first* work that derives the equilibrium bidding strategies for traders in double auctions and analyses the effect of market fees on these strategies. Then we analysed the traders' equilibrium market selection and bidding strategies in the single-home trading environment with multiple marketplaces. Furthermore, we extended the analysis by considering multi-home and hybrid trading environments and different good properties, which is also the *first* work to consider these factors in the analysis of competing double auction marketplaces. We then analysed the effects of different types of fees on obtaining market profits and keeping traders in a single marketplace environment, and showed that the transaction price percentage fee is the most effective in terms of making profits and keeping traders. Finally, we analysed how competing marketplaces set fees in equilibrium. From this analysis, we found that competing marketplaces need to charge high profit fees in equilibrium since traders hide their actual profits by shading.

Finally, in Chapter 5, we analysed the traders' behaviour of selecting marketplaces and submitting shouts in the specific context of the CAT competition. Based on this, we analysed how different market policies affect the performance of competing marketplaces in different environments where different bidding strategies are adopted. This work addressed research challenge 3 of analysing market policies. Finally, we used insights from analysing the market policies and the equilibrium charging strategies to design a competing marketplace. As we have shown, this marketplace performed well in 2010 CAT competition. In particular, it ranked first in the second day's competition and second in the third day's competition. This work addressed research challenge 5 of designing an effective competing marketplace.

When taken together, this research work has successfully addressed the research challenges outlined in the beginning of the thesis and has made a number of important contributions to the design of effective competing double auction marketplaces.

6.2 Future Work

Despite these accomplishments, there still exist limitations of our work. For example, our analysis of market policies is still experimental, and we modelled our game as a one-shot game. Therefore, there are still a number of issues to be addressed in order to further improve the design of competing double auction marketplaces. Specifically, these are as follows:

- In this thesis, we experimentally analysed how different market policies affect the performance of competing marketplaces. However, this analysis was restricted to the specific context of the CAT competition. In the future, we would like to analyse how marketplaces compete with each other in terms of setting their mechanisms (i.e. market policies) from a theoretical perspective. Specifically, we would like to analyse when and how to execute possible transactions (i.e. timing and matching policies) and how to set the transaction prices (i.e. pricing policy). For the timing policy, the market can be cleared when a new shout is admitted (i.e. continuous clearing, in which traders can buy or sell goods quickly, but cannot guarantee traders' profits), or when all traders have submitted their shouts (i.e. round clearing, in which traders' profits can be guaranteed, but they take longer to buy or sell goods). In the real world, in addition to caring about profits made in each marketplace, traders may also look at the time costs of sale or purchase in each marketplace. Different traders may weight differently the obtained profits and time costs, and thus prefer different marketplaces. Therefore, we would like to analyse how to set an effective timing policy to attract traders that have different requirements. For the matching policy, we want to find a matching policy that can guarantee traders' profits (as the equilibrium matching does) and can maximise the number of transactions (as the maximising volume matching does). Furthermore, for the pricing policy, many potential policies could be adopted. This policy determines the trading surplus allocation between buyers and sellers. An improper surplus allocation may cause asymmetric demand and supply, and thus drives traders with excess demand or supply to leave the marketplace. Furthermore, in the competition, other factors may also cause asymmetric demand and supply in the marketplace. For example, the marketplace's opponents may attract more buyers by allocating more profits to them. Faced with this, sellers with excess supply may leave the marketplace since they cannot make transactions. In this situation, the marketplace needs to use the pricing policy to adjust the surplus allocation to re-balance the demand and supply. In the future, we would like to analyse this policy as well. Finally, in addition to theoretically analysing each type of market policy separately, we would like to consider the combination of market policies as a whole, and analyse which one is the most effective. This analysis is likely to provide further and general insights about designing market policies, and can be applied in the more general scenario of competing marketplaces.

- Furthermore, our current theoretical analysis is restricted to one trading round, i.e. a one-shot game. However, the practical competition between marketplaces usually involves multiple trading rounds, and thus is a repeated game. For example, as we mentioned earlier, the CAT competition involves multiple days, each of which includes multiple trading rounds, and thus it is a repeated game. Given this, instead of determining its action in an isolated round, in the repeated game, a trader should take into account the impact of its current action (in terms of selecting marketplaces and submitting shouts) on the future actions of other traders and marketplaces in the future trading rounds, and the same for a competing marketplace. In this vein, we would therefore like to extend our theoretical analysis to a repeated game involving multiple trading rounds. In more detail, in this repeated game, we want to analyse how traders select marketplaces and submit shouts, and how competing marketplaces set market policies and fees. This analysis will provide further insights about designing market policies and charging strategies.

- Finally, although our design of market and charging policies performed well in the 2010 CAT competition, it failed to obtain the first position. This may mean that our design can be further improved. Therefore, we would like to use insights from the above theoretical analysis of market policies and the repeated game to improve the market policy design. Furthermore, we would like to use insights from the above analysis of the repeated game to improve the charging policy. Moreover, as we introduced in Section 5.4, in addition to the gap between the one-shot game and the repeated game, we also made other assumptions, which meant that we could not directly use the equilibrium charging strategies in the CAT competition. We would like to close these gaps in the future. For example, although we cannot guarantee that other marketplaces also use equilibrium charging strategies, according to these marketplaces' fees and the traders' actions, and taking into account the impact of our marketplace's fees on future actions of traders and other marketplaces, we can derive the best response fees to maximise the market profits while maintaining the number of traders at a good level.

Appendix A

Expansion of Equations when Deriving Nash Equilibrium Analytically in Section 3.2.2

Now we expand equations in Section 3.2.2 to derive Nash equilibrium. In Section 3.2.2, we consider 2 buyers, 2 sellers and 2 competing marketplaces only charging profit fees. For trader types, we consider two discrete trader types, rich and poor. Furthermore, we use $\Lambda_1 = t_2^b - t_1^s$ to represent the surplus of a transaction between rich buyer and rich seller, $\Lambda_2 = t_2^b - t_2^s$ to represent the surplus of a transaction between rich buyer and poor seller, $\Lambda_3 = t_1^b - t_1^s$ to represent the surplus of a transaction between poor buyer and rich seller and $\Lambda_4 = t_1^b - t_2^s$ to represent the surplus of a transaction between poor buyer and poor seller. According to Equation 3.14, we obtain the expected utility of rich buyer in marketplace 1:

$$\begin{aligned}
&\tilde{U}_1^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b) \\
&= \Bigg[\Big(1 - \frac{\omega^b(t_2^b, 1, \bar{P}) + \omega^b(t_1^b, 1, \bar{P})}{2}\Big) * \Big(\omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_2^s, 1, \bar{P}) - \frac{\big(\omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_2^s, 1, \bar{P})\big)^2}{4}\Big) * \frac{\Lambda_1 + \Lambda_2}{2} \\
&+ \frac{1}{2} * \Bigg(\omega^b(t_2^b, 1, \bar{P}) * \Big(\omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_2^s, 1, \bar{P}) - \frac{\big(\omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_2^s, 1, \bar{P})\big)^2}{4}\Big) * \frac{\Lambda_1 + \Lambda_2}{2} \\
&+ \omega^b(t_1^b, 1, \bar{P}) * \frac{\big(\omega^s(t_1^s, 1, \bar{P}) + \omega^s(t_2^s, 1, \bar{P})\big)^2}{4} * \frac{\big(\omega^s(t_1^s, 1, \bar{P})\big)^2 * \Lambda_1 + \big(\omega^s(t_2^s, 1, \bar{P})\big)^2 * \Lambda_2}{\big(\omega^s(t_1^s, 1, \bar{P})\big)^2 + \big(\omega^s(t_2^s, 1, \bar{P})\big)^2}\Bigg)\Bigg] * k_1 * (1 - q_1)
\end{aligned} \tag{A.1}$$

and we obtain its expected utility in marketplace 2:

$$\begin{aligned}
&\tilde{U}_2^b(\bar{P}, \bar{K}, \bar{\omega}^b(\bar{P}), \bar{\omega}^s(\bar{P}), t_2^b) \\
&= \Bigg[\frac{\omega^b(t_2^b, 1, \bar{P}) + \omega^b(t_1^b, 1, \bar{P})}{2} * \Big(2 - \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_2^s, 1, \bar{P}) - \frac{\big(2 - \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_2^s, 1, \bar{P})\big)^2}{4}\Big) * \frac{\Lambda_1 + \Lambda_2}{2} \\
&+ \frac{1}{2} * \Bigg(\Big(1 - \omega^b(t_2^b, 1, \bar{P})\Big) * \Big(2 - \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_2^s, 1, \bar{P}) - \frac{\big(2 - \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_2^s, 1, \bar{P})\big)^2}{4}\Big) * \frac{\Lambda_1 + \Lambda_2}{2} \\
&+ \Big(1 - \omega^b(t_1^b, 1, \bar{P})\Big) * \frac{\big(2 - \omega^s(t_1^s, 1, \bar{P}) - \omega^s(t_2^s, 1, \bar{P})\big)^2}{4} * \frac{\big(1 - \omega^s(t_1^s, 1, \bar{P})\big)^2 * \Lambda_1 + \big(1 - \omega^s(t_2^s, 1, \bar{P})\big)^2 * \Lambda_2}{\big(1 - \omega^s(t_1^s, 1, \bar{P})\big)^2 + \big(1 - \omega^s(t_2^s, 1, \bar{P})\big)^2}\Bigg)\Bigg] * k_2 * (1 - q_2)
\end{aligned} \tag{A.2}$$

The poor buyer's expected utility in marketplace 1 is:

$$
\begin{aligned}
&\tilde{U}_1^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^b)\\
&=\Bigg[\Big(1-\frac{\omega^b(t_2^b,1,\bar{P})+\omega^b(t_1^b,1,\bar{P})}{2}\Big)*\Big(\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})-\frac{\big(\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_3+\Lambda_4}{2}\\
&+\frac{1}{2}*\Bigg(\omega^b(t_1^b,1,\bar{P})*\Big(\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})-\frac{\big(\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_3+\Lambda_4}{2}\\
&+\omega^b(t_2^b,1,\bar{P})*\frac{\big(\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})\big)^2}{4}*\frac{\big(\omega^s(t_1^s,1,\bar{P})\big)^2*\Lambda_3+\big(\omega^s(t_2^s,1,\bar{P})\big)^2*\Lambda_4}{\big(\omega^s(t_1^s,1,\bar{P})\big)^2+\big(\omega^s(t_2^s,1,\bar{P})\big)^2}\Bigg)\Bigg]*k_1*(1-q_1)
\end{aligned}
\tag{A.3}
$$

and the poor buyer's expected utility in marketplace 2 is:

$$
\begin{aligned}
&\tilde{U}_2^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^b)\\
&=\Bigg[\frac{\omega^b(t_2^b,1,\bar{P})+\omega^b(t_1^b,1,\bar{P})}{2}*\Big(2-\omega^s(t_1^s,1,\bar{P})-\omega^s(t_2^s,1,\bar{P})-\frac{\big(2-\omega^s(t_1^s,1,\bar{P})-\omega^s(t_2^s,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_3+\Lambda_4}{2}\\
&+\frac{1}{2}*\Bigg(\big(1-\omega^b(t_1^b,1,\bar{P})\big)*\Big(2-\omega^s(t_1^s,1,\bar{P})-\omega^s(t_2^s,1,\bar{P})-\frac{\big(2-\omega^s(t_1^s,1,\bar{P})-\omega^s(t_2^s,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_3+\Lambda_4}{2}\\
&+\big(1-\omega^b(t_2^b,1,\bar{P})\big)*\frac{\big(2-\omega^s(t_1^s,1,\bar{P})-\omega^s(t_2^s,1,\bar{P})\big)^2}{4}*\frac{\big(1-\omega^s(t_1^s,1,\bar{P})\big)^2*\Lambda_3+\big(1-\omega^s(t_2^s,1,\bar{P})\big)^2*\Lambda_4}{\big(1-\omega^s(t_1^s,1,\bar{P})\big)^2+\big(1-\omega^s(t_2^s,1,\bar{P})\big)^2}\Bigg)\Bigg]*k_2*(1-q_1)
\end{aligned}
\tag{A.4}
$$

The rich seller's expected utility in marketplace 1 is:

$$
\begin{aligned}
&\tilde{U}_1^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^s)\\
&=\Bigg[\Big(1-\frac{\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})}{2}\Big)*\Big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})-\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_1+\Lambda_3}{2}\\
&+\frac{1}{2}*\Bigg(\omega^s(t_1^s,1,\bar{P})*\Big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})-\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_1+\Lambda_3}{2}\\
&+\omega^s(t_2^s,1,\bar{P})*\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}*\frac{\big(\omega^b(t_2^b,1,\bar{P})\big)^2*\Lambda_1+\big(\omega^b(t_1^b,1,\bar{P})\big)^2*\Lambda_3}{\big(\omega^b(t_1^b,1,\bar{P})\big)^2+\big(\omega^b(t_2^b,1,\bar{P})\big)^2}\Bigg)*(1-k_1)*(1-q_1)
\end{aligned}
\tag{A.5}
$$

and rich seller's expected utility in marketplace 2 is:

$$
\begin{aligned}
&\tilde{U}_2^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^s)\\
&=\Bigg[\frac{\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})}{2}*\Big(2-\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})-\frac{2-\big(\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_1+\Lambda_3}{2}\\
&+\frac{1}{2}*\Bigg(\big(1-\omega^s(t_1^s,1,\bar{P})\big)*\Big(2-\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})-\frac{\big(2-\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_1+\Lambda_3}{2}\\
&+\big(1-\omega^s(t_2^s,1,\bar{P})\big)*\frac{2-\big(\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}*\frac{\big(1-\omega^b(t_2^b,1,\bar{P})\big)^2*\Lambda_1+\big(1-\omega^b(t_1^b,1,\bar{P})\big)^2*\Lambda_3}{\big(1-\omega^b(t_1^b,1,\bar{P})\big)^2+\big(1-\omega^b(t_2^b,1,\bar{P})\big)^2}\Bigg)*(1-k_2)*(1-q_2)
\end{aligned}
\tag{A.6}
$$

The poor seller's expected utility in marketplace 1 is:

$$\tilde{U}_1^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^s)$$
$$=\Bigg[\Big(1-\frac{\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})}{2}\Big)*\Big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})-\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_2+\Lambda_4}{2}$$
$$+\frac{1}{2}*\Bigg(\omega^s(t_2^s,1,\bar{P})*\Big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})-\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_2+\Lambda_4}{2}$$
$$+\omega^s(t_1^s,1,\bar{P})*\frac{\big(\omega^b(t_1^b,1,\bar{P})+\omega^b(t_2^b,1,\bar{P})\big)^2}{4}*\frac{\big(\omega^b(t_2^b,1,\bar{P})\big)^2*\Lambda_2+\big(\omega^b(t_1^b,1,\bar{P})\big)^2*\Lambda_4}{\big(\omega^b(t_1^b,1,\bar{P})\big)^2+\big(\omega^b(t_2^b,1,\bar{P})\big)^2}\Bigg)*(1-k_1)*(1-q_1) \quad (A.7)$$

and the poor seller's expected utility in marketplace 2 is:

$$\tilde{U}_2^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^s)$$
$$=\Bigg[\frac{\omega^s(t_1^s,1,\bar{P})+\omega^s(t_2^s,1,\bar{P})}{2}*\Big(2-\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})-\frac{2-\big(\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_2+\Lambda_4}{2}$$
$$+\frac{1}{2}*\Bigg(\Big(1-\omega^s(t_2^s,1,\bar{P})\Big)*\Big(2-\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})-\frac{2-\big(\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}\Big)*\frac{\Lambda_2+\Lambda_4}{2}$$
$$+\Big(1-\omega^s(t_1^s,1,\bar{P})\Big)*\frac{2-\big(\omega^b(t_1^b,1,\bar{P})-\omega^b(t_2^b,1,\bar{P})\big)^2}{4}*\frac{\big(1-\omega^b(t_2^b,1,\bar{P})\big)^2*\Lambda_2+\big(1-\omega^b(t_1^b,1,\bar{P})\big)^2*\Lambda_4}{\big(1-\omega^b(t_1^b,1,\bar{P})\big)^2+\big(1-\omega^b(t_2^b,1,\bar{P})\big)^2}\Bigg)\Bigg]*(1-k_2)*(1-q_2) \quad (A.8)$$

Recall that in the mixed Nash equilibrium, we have:

$$\tilde{U}_1^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^b) = \tilde{U}_2^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^b) \quad (A.9)$$
$$\tilde{U}_1^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^b) = \tilde{U}_2^b(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^b) \quad (A.10)$$
$$\tilde{U}_1^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^s) = \tilde{U}_2^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_1^s) \quad (A.11)$$
$$\tilde{U}_1^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^s) = \tilde{U}_2^s(\bar{P},\bar{K},\bar{\omega}^b(\bar{P}),\bar{\omega}^s(\bar{P}),t_2^s) \quad (A.12)$$

Now we replace the left-hand and right-hand sides of Equations A.9, A.10, A.11 and A.12 by Equations A.1, A.2, A.3, A.4, A.5, A.6, A.7 and A.8 respectively. Then when $\omega^b(t_1^b,1,\bar{P}) = \omega^b(t_2^b,1,\bar{P})$ and $\omega^s(t_1^s,1,\bar{P}) = \omega^s(t_2^s,1,\bar{P})$, Equations A.9 and A.10 can be rewritten as the same equation:

$$\Big[2*\omega^s(t_1^s,1,\bar{P})-\big(\omega^s(t_1^s,1,\bar{P})\big)^2-\omega^b(t_1^b,1,\bar{P})*\omega^s(t_1^s,1,\bar{P})+\omega^b(t_1^b,1,\bar{P})*\big(\omega^s(t_1^s,1,\bar{P})\big)^2\Big]*k_1*(1-q_1)$$
$$=\Big[1-\omega^s(t_1^s,1,\bar{P})+\omega^b(t_1^b,1,\bar{P})*\omega^s(t_1^s,1,\bar{P})-\omega^b(t_1^b,1,\bar{P})*\big(\omega^s(t_1^s,1,\bar{P})\big)^2\Big]*k_2*(1-q_2) \quad (A.13)$$

and Equations A.11 and A.12 can be rewritten as the same equation:

$$\Big[2*\omega^b(t_1^b,1,\bar{P})-\big(\omega^b(t_1^b,1,\bar{P})\big)^2-\omega^b(t_1^b,1,\bar{P})*\omega^s(t_1^s,1,\bar{P})+\omega^s(t_1^s,1,\bar{P})*\big(\omega^b(t_1^b,1,\bar{P})\big)^2\Big]*(1-k_1)*(1-q_1)$$
$$=\Big[1-\omega^b(t_1^b,1,\bar{P})+\omega^b(t_1^b,1,\bar{P})*\omega^s(t_1^s,1,\bar{P})-\omega^s(t_1^s,1,\bar{P})*\big(\omega^b(t_1^b,1,\bar{P})\big)^2\Big]*(1-k_2)*(1-q_2) \quad (A.14)$$

This means that $\omega^b(t_1^b,1,\bar{P}) = \omega^b(t_2^b,1,\bar{P})$ and $\omega^s(t_1^s,1,\bar{P}) = \omega^s(t_2^s,1,\bar{P})$ will be one of the solutions for Equations A.9, A.10, A.11 and A.12, and by solving Equations A.13 and A.14, we can obtain the solutions.

Appendix B

An Alternative Approach to Calculate Expected Utilities of Traders in Section 4.1.2

In Section 4.1.2, when we calculate the expected utilities of traders, we need to consider all possible numbers of buyers and sellers choosing different actions. As we discussed previously, this calculation is demanding. Given this fact, we introduce an alternative approach to calculate a trader's expected utility. Specifically, we calculate the expected utility of a buyer with type θ^b adopting the action $\delta^b = \langle d_1^b, d_2^b, ..., d_M^b \rangle$ given the other buyers' action distribution Ω^b and the sellers' action distribution Ω^s, and the fee system $\bar{P}$. The calculation for the seller is analogous. The expected utility consists of two parts: the *expected value* on the goods and the *expected payment*. In the following, we derive these two parts respectively.

We first derive the buyer's expected value on traded goods. In order to do this, we need to know the buyer's joint positions. In contrast to that in Section 4.1.2, we derived the buyer's joint positions from the number of buyers choosing different actions, we derive this by comparing buyers' actions in terms of comparing bids in these actions. In more detail, when comparing the buyer's bid d_m^b with another bid $d_m'^b$ in marketplace m, there are three possible mutually exclusive events, i.e. $d_m^b > d_m'^b$ (i.e. the buyer's bid in marketplace m is higher than another buyer's bid in this marketplace)[1], $d_m^b = d_m'^b$ (i.e. both buyers have the same bid in marketplace m) or $d_m^b < d_m'^b$ (i.e. the buyer's bid in marketplace m is less than another buyer's bid in this marketplace). Thus when comparing the buyer's action, $\langle d_1^b, d_2^b, ..., d_M^b \rangle$, with another buyer's action, $\langle d_1'^b, d_2'^b, ..., d_M'^b \rangle$, in terms of comparing bids in each marketplace, there are 3^M possible mutually exclusive events, each of which is a *joint event* of comparing bids across M marketplaces. We use $\mathcal{R} = \{R_1, R_2, ..., R_{3^M}\}$ to represent the set of all these possible joint events. For example, when there are two marketplaces, we have $\mathcal{R} = \{R_1 = (<,<), R_2 = (<,=), R_3 = (<,>), R_4 = (=,<), R_5 = (=,=), R_6 = (=,>), R_7 = (>,<), R_8 = (>,=), R_9 = (>,>)\}$. As an example, event

[1]For not choosing the marketplace, i.e. $d_m'^b = \ominus$, we regard that $d_m^b > d_m'^b = \ominus$.

$R_1 = (<,<)$ means that the buyer's bids in both marketplaces are less than the bids placed by another buyer. The probability that a joint event R_i occurs is:

$$\phi_i^b = \sum_{\delta_j^b \in \Delta : \delta^b \; ? \; \delta_j^b = R_i} \omega_j^b \tag{B.1}$$

where as we introduced previously, ω_j^b is the probability of a buyer choosing the action δ_j^b, and $\delta^b \; ? \; \delta_j^b = R_i$ means that event R_i occurs when we compare action δ^b with δ_j^b.

To calculate the position, we furthermore need to know the number of buyers satisfying each event. Specifically, we use a 3^M-tuple $\bar{x} = \langle x_1, ..., x_{3^M} \rangle \in \mathcal{X}$ to represent the number of buyers satisfying each event, where there are exactly x_i buyers satisfying event R_i when they compare their actions with the buyer's action, and $\mathcal{X}$ is the set of all such possible tuples satisfying the conditions $x_i \geq 0$ and $\sum_{i=1}^{3^M} x_i = B - 1$ (note that we need to exclude the buyer for which we are calculating the expected utility) and $|\mathcal{X}| = 3^{M*(B-1)}$. The probability of exactly x_i buyers satisfying event R_i is $(\phi_i^b)^{x_i}$, and then the probability of this tuple appearing is:

$$\rho^b(\bar{x}) = \binom{B-1}{x_1, ..., x_{3^M}} * \prod_{i=1}^{3^M} (\phi_i^b)^{x_i} \tag{B.2}$$

Now given tuple $\bar{x}$, we determine the buyer's joint positions as follows. Firstly, we obtain the number of buyers whose bids are greater than the buyer's bid in marketplace m, which is given by:

$$X_m^>(\bar{x}) = \sum_{R_i \in \mathcal{R} : R_{im} = \text{`}<\text{'}} x_i \tag{B.3}$$

where R_{im} is the event of comparing bids in marketplace m from the joint event R_i. Similarly, we use $X_m^=(\bar{x})$ to represent the number of buyers whose bids are equal to the buyer's bid in marketplace m (excluding the buyer itself):

$$X_m^=(\bar{x}) = \sum_{R_i \in \mathcal{R} : R_{im} = \text{`}=\text{'}} x_i \tag{B.4}$$

Due to having discrete bids and given $X_m^>(\bar{x})$ buyers bidding higher than the buyer's bid d_m^b and $X_m^=(\bar{x})$ buyers bidding equal to d_m^b, the buyer's position in marketplace m could be anywhere from $X_m^>(\bar{x}) + 1$ to $X_m^>(\bar{x}) + X_m^=(\bar{x}) + 1$, which constitutes the buyer's position range in this marketplace. Since $X_m^=(\bar{x}) + 1$ buyers have the same bid, a tie-breaking rule is needed to determine the buyer's position. As we did in Section 4.1.2, we adopt a standard tie-breaking rule where each of these possible positions[2] occurs with equal probability, i.e. $1/(X_m^=(\bar{x}) + 1)$. For example, when there are two marketplaces and joint event R_i occurs for x_i buyers, $i = 1, ..., 9$, in marketplace 1, we have $X_1^>(\bar{x}) = x_1 + x_2 + x_3$ and $X_1^=(\bar{x}) = x_4 + x_5 + x_6$, and the buyer's position will be anywhere from $X_1^>(\bar{x}) + 1$ to $X_1^>(\bar{x}) + X_1^=(\bar{x}) + 1$ with equal probability $1/(X_1^=(\bar{x}) + 1)$; and in marketplace 2, we have $X_2^>(\bar{x}) = x_1 + x_4 + x_7$ and $X_2^=(\bar{x}) = x_2 + x_5 + x_8$, and the buyer's position will be anywhere from $X_2^>(\bar{x}) + 1$ to $X_2^>(\bar{x}) + X_2^=(\bar{x}) + 1$ with equal probability $1/(X_2^=(\bar{x}) + 1)$. Now, given

[2] They are $X_m^>(\bar{x}) + 1$, $X_m^>(\bar{x}) + 2$,..., $X_m^>(\bar{x}) + X_m^=(\bar{x}) + 1$.

the buyer's position ranges in different marketplaces, we can obtain the set of all possible joint positions for the buyer. Specifically, we use a M-tuple $\bar{v}_{\bar{x}} = \langle v_1, ..., v_M \rangle \in \mathcal{V}_{\bar{x}}$ to represent one of the possible joint positions where 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 $X_m^{>}(\bar{x}) + 1 \leq v_m \leq X_m^{>}(\bar{x}) + X_m^{=}(\bar{x}) + 1$ $(m = 1, ..., M)$. The probability of the buyer having the joint positions $\bar{v}_{\bar{x}}$ given the tuple $\bar{x}$ is:

$$\Phi(\bar{v}_{\bar{x}}) = \prod_{m=1}^{M} \frac{1}{X_m^{=}(\bar{x}) + 1} \tag{B.5}$$

Note that tie-breaking occurs independently for each marketplace.

In addition to depending on positions in different marketplaces, the buyer's expected value also depends on the numbers of sellers choosing different actions. Different from Section 4.1.2 where we consider all possible numbers of sellers choosing different actions, we consider the number of sellers submitting asks in each marketplace. Given this and the sellers' action distributions, we can know whether the buyer will be matched in each marketplace. Since a seller can place multiple asks in multiple marketplaces at the same time, the numbers of sellers submitting asks in different marketplaces are also correlated with each other. Thus we need to consider the joint numbers of sellers submitting asks in different marketplaces. Specifically, we use a M-tuple $\bar{y} = \langle y_1, ..., y_M \rangle \in \mathcal{Y}$ to denote the joint numbers of sellers submitting asks in different marketplaces, where y_m is the number of sellers submitting asks in marketplace m, and $\mathcal{Y}$ is the set of all such possible tuples satisfying the condition that $0 \leq y_m \leq S$ $(m = 1, ..., M)$, and $|\mathcal{Y}| = (S + 1)^M$. In the following, we derive the probability of a specific tuple $\bar{y}$ appearing. We use the powerset $2^{\mathcal{M}} = \{\mathcal{M}_1, ..., \mathcal{M}_{2^M}\}$ to denote the set of all possible marketplace subsets. For marketplace subset $\mathcal{M}_I$, $I = 1, ..., 2^M$, we use $\Delta_I \subset \Delta$ to denote the subset of seller's actions which only submit asks in $\mathcal{M}_I$, and submit $\ominus$ in the complement $\mathcal{M} - \mathcal{M}_I$ (i.e. sellers do not choose marketplaces in $\mathcal{M} - \mathcal{M}_I$). From the seller's action distribution Ω^s, we can calculate the probability of sellers choosing actions from Δ_I (i.e. only submitting asks in marketplace subset $\mathcal{M}_I$):

$$\mu_I^s = \sum_{\delta_j^s \in \Delta_I} \omega_j^s \tag{B.6}$$

Now we use a 2^M-tuple $\bar{z} = \langle z_1, ..., z_{2^M} \rangle \in \mathcal{Z}$, $\sum_{I=1}^{2^M} z_I = S$, to denote the numbers of sellers submitting asks in each of these possible marketplace subsets, where z_I is the number of sellers choosing actions from action subset Δ_I (i.e. only submitting asks in subset $\mathcal{M}_I$), and $\mathcal{Z}$ is the set of all such possible tuples. The probability of this tuple appearing is:

$$\rho^s(\bar{z}) = \binom{S}{z_1, ..., z_{2^M}} * \prod_{I=1}^{2^M} (\mu_I^s)^{z_I} \tag{B.7}$$

Given tuple $\bar{z}$, we can know the joint numbers of sellers submitting asks in different marketplaces. Specifically, we introduce a function $g(\bar{z}) = \langle y'_1, ..., y'_M \rangle$, which converts the numbers of sellers submitting asks in different marketplace subsets to the numbers of sellers submitting asks in different single marketplaces, where

$$y'_m = \sum_{\mathcal{M}_I \in 2^{\mathcal{M}}: m \in \mathcal{M}_I} z_I \tag{B.8}$$

Now we calculate the probability of the tuple $\bar{y}$ appearing:

$$\rho^s(\bar{y}) = \sum_{\bar{z} \in \mathcal{Z}: \bar{y} = g(\bar{z})} \rho^s(\bar{z}) \tag{B.9}$$

Now given the buyer's joint positions $\bar{v}_{\bar{x}}$, the joint numbers of sellers submitting asks in different marketplaces $\bar{y}$, we are ready to calculate its expected value on traded goods. Since the buyer can enter multiple marketplaces and thus purchase multiple goods, we need to consider its expected value on different units of goods. Remember that 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 4.1.1, we have defined the buyer's value $v^b(\theta^b, T)$ on T units of goods by considering different good properties (see Equation 4.1). Now by considering all possible marketplace subsets with cardinality T, where exactly T transactions are made by this buyer, we obtain the buyer's expected value when it purchases T units of goods given its joint positions $\bar{v}_{\bar{x}}$ and the numbers of sellers submitting asks in different marketplaces $\bar{y}$:

$$\begin{aligned} \tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) &= \sum_{\mathcal{M}_I \subset 2^{\mathcal{M}}: |\mathcal{M}_I| = T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) * v^b(\theta^b, T) \\ &= \sum_{\mathcal{M}_I \subset 2^{\mathcal{M}}: |\mathcal{M}_I| = T} \varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) * \alpha^b_T * \theta^b \end{aligned} \tag{B.10}$$

where $\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I)$ is the probability of the buyer making transactions in marketplaces $\mathcal{M}_I$ and not making transactions in $\mathcal{M} - \mathcal{M}_I$. Note that given the buyer's position and the number of sellers submitting asks in each marketplace, the probability of the buyer making a transaction (or not making a transaction) in each marketplace is independent of each other, and thus the probability of the buyer making transactions in $\mathcal{M}_I$ and not making transactions in $\mathcal{M} - \mathcal{M}_I$ is given by:

$$\varphi^b(\bar{v}_{\bar{x}}, \bar{y}, \delta^b, \mathcal{M}_I) = \sum_{m \in \mathcal{M}_I} \psi^b(v_m, m, d^b_m, y_m) * \sum_{m \in \mathcal{M} - \mathcal{M}_I} \chi^b(v_m, m, d^b_m, y_m) \tag{B.11}$$

where $\psi^b(v_m, m, d^b_m, y_m)$ is the probability of the buyer with bid d^b_m making a transaction in marketplace m given its position v_m and y_m sellers submitting asks in this marketplace, and $\chi^b(v_m, m, d^b_m, y_m)$ is the probability of the buyer with bid d^b_m not making a transaction in marketplace m. In the following, we derive them respectively.

First, we introduce three support functions[3]: $e^<_s(d)$ denotes the probability that the sellers' asks

[3] These are calculated by taking the sum of the probabilities of actions whose corresponding shouts in marketplace m satisfy the conditions defined by these functions.

are strictly less than d in marketplace m, $e_s^{=}(d)$ denotes the probability that the sellers' asks are equal to d in marketplace m, and $e_s^{>}(d)$ denotes the probability that the sellers' asks are strictly higher than d in marketplace m but are not $\ominus$. Given the buyer's position v_m, its bid d_m^b, and the number of sellers y_m submitting asks in marketplace m, the probability of the buyer making a transaction in marketplace m is given by:

$$\psi^b(v_m, m, d_m^b, y_m) = \sum_{c=v_m}^{y_m} \binom{y_m}{c} * \left(\frac{e_s^{<}(d_m^b) + e_s^{=}(d_m^b)}{E}\right)^c * \left(\frac{e_s^{>}(d_m^b)}{E}\right)^{y_m - c} \quad \text{(B.12)}$$

where $E = e_s^{<}(d_m^b) + e_s^{=}(d_m^b) + e_s^{>}(d_m^b)$. Note that this calculation is given the condition that y_m sellers have submitted asks in marketplace m, and thus the probability of an ask less than or equal to d_m^b in marketplace m should be normalised, i.e. $(e_s^{<}(d_m^b) + e_s^{=}(d_m^b))/E$, and the same reason for $e_s^{>}(d_m^b)/E$. The probability of the buyer not making a transaction in marketplace m is given by:

$$\chi^b(v_m, m, d_m^b, y_m) = \sum_{c=0}^{v_m - 1} \binom{y_m}{c} * \left(\frac{e_s^{<}(d_m^b) + e_s^{=}(d_m^b)}{E}\right)^c * \left(\frac{e_s^{>}(d_m^b)}{E}\right)^{y_m - c} \quad \text{(B.13)}$$

Finally, by considering all possible numbers of units the buyer purchase, all possible joint numbers of sellers submitting asks in different marketplaces, all possible joint positions and all possible numbers of buyers satisfying different joint events, the buyer's expected value is given by:

$$\tilde{V}(\theta^b, \delta^b, \Omega^b, \Omega^s) = \sum_{\bar{x} \in \mathcal{X}} \rho^b(\bar{x}) * \sum_{\bar{v}_{\bar{x}} \in \mathcal{V}_{\bar{x}}} \upsilon(\bar{v}_{\bar{x}}) * \sum_{\bar{y} \in \mathcal{Y}} \rho^s(\bar{y}) * \sum_{T=1}^{M} \tilde{V}(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, T) \quad \text{(B.14)}$$

After deriving the expected value, in the following, we derive the expected payment of the buyer given the action distributions of buyers and sellers, Ω^b and Ω^s, and the fee system $\bar{P}$. Firstly, we derive the buyer's expected payment given its joint positions $\bar{v}_{\bar{x}}$ and joint numbers of sellers submitting asks in different marketplaces $\bar{y}$. The buyer's expected payment is the sum of its expected payment in each marketplace, which is:

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

where $\tilde{\mathcal{P}}_m^b(v_m, y_m, \theta^b, d_m^b, \Omega^b, \Omega^s, p_m)$ is the buyer's expected payment in marketplace m when it bids d_m^b given its position v_m and y_m sellers submitting asks in this marketplace, and it is given by:

$$\tilde{\mathcal{P}}_m^b(v_m, y_m, \theta^b, d_m^b, \Omega^b, \Omega^s, p_m) = \begin{cases} 0 & \textbf{\textit{if }} d_m^b = \ominus \\ \sum_{d^s \in \Phi - \{\ominus\}} \tilde{\mathcal{P}}_m^b(v_m, y_m, \theta^b, d_m^b, \Omega^b, \Omega^s, d^s, p_m) + r_m + \varepsilon & \textbf{\textit{if }} d_m^b \neq \ominus \end{cases}$$

where $\sum_{d^s\in\Phi-\{\ominus\}} \tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, d^s, p_m)$ is the buyer's expected payment excluding registration fee r_m and constant cost ε, and $\tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, d^s, p_m)$ is the buyer's expected payment when it attempts to be matched with the ask d^s, which is given by:

$$\tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, d^s, p_m) = \begin{cases} 0 & \textbf{\textit{if}}\ d^b_m \leq d^s \\ \sum_{y^s_1=0}^{v_m-1} \sum_{y^s_2=v_m-y^s_1}^{y_m-y^s_1} \rho^s(y_m, y^s_1, y^s_2, d^s) * \tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, p_m|d^s) & \textbf{\textit{if}}\ d^b_m > d^s \end{cases}$$

where

$$\rho^s(y_m, y^s_1, y^s_2, d^s) = \binom{y_m}{y^s_1, y^s_2, y_m - y^s_1 - y^s_2} * \prod_{i=1}^{3} \left(\frac{e^s_i(d^s)}{E}\right)^{y^s_i}$$

is the probability that there are exactly y^s_1 asks strictly less than d^s, exactly y^s_2 asks equal to d^s (including the ask itself), and exactly $y^m - y^s_1 - y^s_2$ asks greater than d^s. The same as before, since the calculation is given the condition that y_m sellers have submitted asks in marketplace m, and thus the probability of an ask less than (equal to, or greater than) d^s in marketplace m should be normalised, i.e. $e^s_i(d^s)/E$ and $E = e^<_s(d^s) + e^=_s(d^s) + e^>_s(d^s)$. Note that $\sum_{y^s_1=0}^{v_m-1} \sum_{y^s_2=v_m-y^s_1}^{y_m-y^s_1} \rho^s(y_m, y^s_1, y^s_2, d^s)$ actually gives the overall probability that this match happens. Finally, $\tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, p_m|d^s)$ is the buyer's expected payment when it is matched with the ask d^s. This is given by:

$$\tilde{\mathcal{P}}^b_m(v_m, y_m, \theta^b, d^b_m, \Omega^b, \Omega^s, p_m|d^s) = \mathrm{TP} + t_m + \mathrm{TP} * o_m + (d^b_m - \mathrm{TP}) * q_m \tag{B.15}$$

where $\mathrm{TP} = d^s * k_m + d^b_m * (1 - k_m)$ is the transaction price, t_m is the transaction fee, $\mathrm{TP} * o_m$ is the payment of transaction price percentage fee, and $(d^b_m - \mathrm{TP}) * q_m$ is the payment of profit fee.

Now by considering all possible joint numbers of sellers submitting asks in different marketplaces, all possible joint positions and all possible numbers of buyers satisfying different joint events, the buyer's expected payment is given by:

$$\tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) = \sum_{\bar{x}\in\mathcal{X}} \rho^b(\bar{x}) * \sum_{\bar{v}_{\bar{x}}\in\mathcal{V}_{\bar{x}}} \upsilon(\bar{v}_{\bar{x}}) * \sum_{\bar{y}\in\mathcal{Y}} \rho^s(\bar{y}) * \tilde{\mathcal{P}}^b(\bar{v}_{\bar{x}}, \bar{y}, \theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) \tag{B.16}$$

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

$$\tilde{U}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) = \tilde{V}^b(\theta^b, \delta^b, \Omega^b, \Omega^s) - \tilde{\mathcal{P}}^b(\theta^b, \delta^b, \Omega^b, \Omega^s, \bar{P}) \tag{B.17}$$

In this way of calculating a buyer's expected utility, for buyers, by comparing actions in terms of comparing bids in the actions, we reduce the possibilities of buyers' action choices from $|\Phi|^{M*(B-1)}$ to $3^{M*(B-1)}$, and for sellers, by considering the possibilities of the number of the sellers submitting asks in different marketplaces, we reduce the possibilities from $|\Phi|^{M*S}$ to $(S+1)^M$. In our analysis where we consider 11 possible shouts plus $\ominus$, 2 competing marketplaces, 5 buyers and 5 sellers, we significantly reduce the possibilities of buyers' action choices from $12^8 = 429981696$ to $3^8 = 6561$, and reduce the possibilities for sellers from $12^{10} = 61917364224$ to $6^2 = 36$.

Bibliography

P. Anthony and N. R. Jennings. Evolving bidding strategies for multiple auctions. In *Proceedings of the 15th European Conference on Artifical Intelligence*, pages 178–182, 2002.

P. Anthony and N. R. Jennings. Developing a bidding agent for multiple heterogeneous auctions. *ACM Transactions on Internet Technology*, 3:185–217, 2006.

D. Begg, S. Fischer, and R. Dornbusch. *Economics*. McGraw-Hill, 1994.

K. Binmore. *Fun and Games, A Text on Game Theory*. D.C. Heath, 1991.

N. Borissov. Q-strategy: Automated bidding and convergence in computational markets. In *Proceedings of the 21st Innovative Applications of Artificial Intelligence Conference*, pages 54–59, 2009.

N. Borissov, D. Neumann, and C. Weinhardt. Automated bidding in computational markets: An application in market-based allocation of computing services. *Autonomous Agents and Multiagent Systems*, 21:115–142, 2010.

G. W. Brown. Iterative solutions of games by fictitious play. *Activity Analysis of Production and Allocation*, pages 374–376, 1951.

S. Bullock. *Evolutionary simulation models: On their character, and application to problems concerning the evolution of natural signalling systems*. PhD thesis, University of Sussex, 1997.

R. Burguet and J. Sakovics. Imperfect competition in auction designs. *International Economic Review*, 40(1):231–247, 1999.

A. Byde. A comparison among bidding algorithms for multiple auctions. Technical report, Hewlett Packard Research Labs, 2001a.

A. Byde. An optimal dynamic programming model for algorithm in simultaneous auctions. Technical report, Hewlett Packard Research Labs, 2001b.

A. Byde, C. Preist, and N. R. Jennings. Decision procedures for multiple auctions. In *Proceedings of the 1st International Conference on Autonomous Agents and Multi-Agent Systems*, pages 613–620, 2002.

K. Cai, E. H. Gerding, P. McBurney, J. Niu, S. Parsons, and S. Phelps. Overview of CAT: A market design competition (version 2.0). Technical report, Dept. of Computer Science, University of Liverpool, 2009.

K. Cai, J. Niu, and S. Parsons. On the economic effects of competition between double auction markets. In *Proceedings of the 10th Workshop on Agent-Mediated Electronic Commerce*, 2008.

B. Caillaud and B. Jullien. Chicken & egg: Competition among intermediation service providers. *The RAND Journal of Economics*, 34(2):309–328, 2003.

C. Carrie. Illuminating the new dark influence on trading and U.S. market structure. *Journal of Trading*, 3(1):40–55, 2008.

D. Cliff and J. Bruten. Minimal-intelligence agents for bargaining behaviors in market-based environments. Technical report, Hewlett Packard Research Labs, 1997.

E. Damiano and L. Hao. Competing matchmaking. *Journal of European Economic Association*, 6(4):789–818, 2008.

G. Ellison, M. Mobius, and D. Fudenberg. Competing auctions. *Journal of European Economic Association*, 2(1):30–66, 2004.

D. Friedman and J. Rust. *The double auction market: institutions, theories and evidence*, volume XIV of Santa Fe Institute Studies in the Science of Complexity. Perseus Publishing, 1993.

D. Fudenberg and J. Tirole. *Game Theory*. The MIT Press, 1991.

E. H. Gerding, R. Dash, A. Byde, and N. R. Jennings. Optimal strategies for bidding agents participating in simultaneous Vickrey auctions with perfect substitutes. *Journal of Artifical Intelligence Research*, 32:939–982, 2008.

E. H. Gerding, A. Rogers, R. K. Dash, and N. R. Jennings. Sellers competing for buyers in online markets: reserve prices, shill bids, and auction fees. In *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, pages 1287–1293, 2007.

S. Gjerstad and J. Dickhaut. Price formation in double auctions. *Games and Economic Behavior*, 22:1–29, 1998.

D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1): 119–137, 1993.

J. C. Harsanyi. Bargaining in ignorance of the opponent's utility function. *Econometrica*, 18: 155–162, 1962.

M. He and N. R. Jennings. Southamptontac: An adaptive autonomous trading agent. *ACM Transactions on Internet Technology*, 3:218–235, 2003.

M. He, N. R. Jennings, and A. Prugel-Bennett. A heuristic bidding strategy for buying multiple goods in multiple English auctions. *ACM Transactions on Internet Technology*, 6(4):465–496, 2006.

M. He, H. F. Leung, and N. R. Jennings. A fuzzy logic based bidding strategy for autonomous agents in continuous double auctions. *IEEE Transactions on Knowledge and Data Engineering*, 15(6):1345–1363, 2003.

A. Hernando-Veciana. Competition among auctioneers in large markets. *Journal of Economic Theory*, 121(1):107–127, 2005.

S. Honari, M. Ebadi, A. Foshati, and M. Gomrokchi. Price estimation of PersianCAT market equilibrium. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2009.

D. L. Jones. *Auction market theory: a primer*. CISCO, 1988.

V. Krishna. *Auction Theory*. The Academic Press, 2002.

R. S. Lee. Competing platforms. *Working paper*, 2008.

C. E. Lemke and J. T. Howson. Equilibrium points of bimatrix games. *Journal of the Society for Industrial and Applied Mathematics*, 12(2):413–423, 1964.

K. Leyton-Brown and Y. Shoh. *Essentials of Game Theory: A Concise, Multidisciplinary Introduction*. Morgan & Claypool Publishers, 2008.

A. Mas-Collel, W. Whinston, and J. Green. *Microeconomic theory*. Oxford University Press, 1995.

J. S. Maynard. *Evolution and the Theory of Games*. Cambridge University Press, 1982.

R. P. McAfee. Mechanism design by competing sellers. *Econometrica*, 61(6):1281–1312, 1993.

M. H. Miller, J. D. Hawke, B. Malkiel, and M. Scholes. Findings of the committee of inquiry examining the events surrounding October 19, 1987. Technical report, Chicago Mercantile Exchange, 1988.

J. Nicolaisen, V. Petrov, and L. Tesfatsion. Market power and efficiency in a computational electricity market with discriminatory double auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523, 2001.

J. Niu, K. Cai, E. H. Gerding, P. McBurney, and S. Parsons. Characterizing effective auction mechanisms: Insights from the 2007 TAC Market Design Competition. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1079–1086, 2008a.

J. Niu, K. Cai, S. Parsons, E. H. Gerding, P. McBurney, T. Moyaux, S. Phelps, and D. Shield. JCAT: A platform for the TAC Market Design Competition. In *Proceedings of the 7th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1649–1650, 2008b.

J. Niu, K. Cai, S. Parsons, P. McBurney, and E. H. Gerding. What the 2007 tac market design game tells us about effective auction mechanisms. *Autonomous Agents and Multiagent Systems*, 21:172–203, 2010.

J. Niu, K. Cai, S. Parsons, and E. Sklar. Reducing price fluctuation in continuous double auction market through pricing policy and shout improvement. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1143–1150, 2006.

J. Niu, K. Cai, S. Parsons, and E. Sklar. Some preliminary results on competition between markets for automated traders. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2007.

J. Niu, A. Mmoloke, P. McBurney, and S. Parsons. CATP: A communication protocol for CAT games (version 2.0). Technical report, Dept. of Computer Science, University of Liverpool, 2009.

M. J. Osborne and A. Rubinstein. *A Course in Game Theory*. The MIT Press, 1994.

A. Petric, V. Podobnik, A. Grguric, and M. Zemljic. Designing an effective e-market: an overview of the cat agent. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2008.

S. Phelps. *Evolutionary mechanism design*. PhD thesis, University of Liverpool, 2008.

S. Phelps, P. McBurney, and S. Parsons. Evolutionary mechanism design: a review. *Autonomous Agents and Multi-Agent Systems*, 21(2):237–264, 2010.

S. Phelps, P. McBurney, S. Parsons, and E. Sklar. Applying genetic programming to economic mechanism design: Evolving a pricing rule for a continuous double auction. In *Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1096–1097, 2003.

S. Phelps, S. Parsons, and P. McBurney. An evolutionary game-theoretic comparison of two double auction market designs. In *Proceedings of the 6th Workshop on Agent-Mediated Electronic Commerce*, pages 101–114, 2006.

C. Preist, C. Bartolini, and I. Phillips. Algorithm design for agents which participate in multiple simultaneous auctions. In *Proceedings of the 3rd Workshop on Agent-Mediated Electronic Commerce*, pages 139–154, 2001a.

C. Preist, A. Byde, and C. Bartolini. Economic dynamics of agents in multiple auctions. In *Proceedings of the 5th International Conference on Autonomous Agents*, pages 545–551, 2001b.

C. Preist and M. van Tol. Adaptive agents in a persistent shout double auction. Technical report, Hewlett Packard Research Labs, 2003.

Z. Rabinovich, E. H. Gerding, M. Polukarov, and N. R. Jennings. Generalised fictitious play for a continuum of anonymous players. In *Proceedings of the 21st Joint Conference on Artificial Intelligence*, pages 245–250, 2009.

R. Radner and R. W. Rosenthal. Private information and pure-strategy equilibria. *Mathematics of Operations Research*, 7(3):401–409, 1982.

E. R. Robinson, P. McBurney, and X. Yao. How specialised are specialists? Generalisation properties of entries from the 2008 TAC Market Design Competition. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2009.

J. Rochet and J. Tirole. Platform competition in two-sided markets. *Journal of European Economic Association*, 1(4):990–1029, 2003.

A. E. Roth and I. Erev. Predicting how people play games: Reinforcement learning in experimental games with unique, mixed strategy equilibria. *The American Economic Review*, 88(4): 848–881, 1998.

R. Savani and B. Stengel. Hard-to-solve bimatrix games. *Econometrica*, 74(2):397–429, 2006.

A. Shah and S. Thomas. David and goliath: Displacing a primary market. *Global Financial Markets*, 1(1):14–21, 2000.

V. L. Smith. An experimental study of competitive market behavior. *Journal of Political Economy*, 70:111–137, 1962.

J. Sohn, A. Kwasnica, and T. Mullen. Human traders across multiple markets: attracting intramarginal traders under economic experiments. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2010.

J. Sohn, S. Lee, and T. Mullen. Attracting intra-marginal traders across multiple markets. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2009.

P. G. Spirakis. Approximate equilibria for strategic two person games. In *Proceedings of the 1st International Symposium on Algorithmic Game Theory*, pages 5–21, 2008.

L. C. Stavrogiannis and P. A. Mitkas. CAT 2008 post-tournament evaluation: the Mertacor's perspective. In *Proceedings of the Workshop on Trading Agent Design and Analysis*, 2009.

P. Stone, R. E. Schapire, M. L. Littman, J. A. Csirik, and D. McAllester. Decision-theoretic bidding based on learned density models in simultaneous, interacting auctions. *Journal of Artificial Intelligence Research*, 19(1):209–242, 2003.

R. S. Sutton and A. G. Barto. *Reinforcement learning: An introduction*. The MIT Press, 1998.

G. Tesauro and J. L. Bredin. Strategic sequential bidding in auctions using dynamic programming. In *Proceedings of the 1st International Conference on Autonomous Agents and Multi-Agent Systems*, pages 591–598, 2002.

P. Toulis, D. Kehagias, and P. A. Mitkas. Mertacor: a successful autonomous trading agent. In *Proceedings of the 5th International Conference on Autonomous Agents and Multi-Agent Systems*, pages 1191–1198, 2006.

K. Tuyls. *Learning in multi-agent systems, an evolutionary game theoretic approach*. PhD thesis, Vrije Universiteit Brussel, 2004.

I. A. Vetsikas and B. Selman. A principled study of the design tradeoffs for autonomous trading agents. In *Proceedings of the 2nd International Conference on Autonomous Agents and Multi-Agent Systems*, pages 473–480, 2003.

J. von Neumann and G. W. Brown. Solutions of games by differential equations. *Contributions to the Theory of Games*, pages 73–79, 1950.

P. Vytelingum, D. Cliff, and N. R. Jennings. Strategic bidding in continuous double auctions. *Artificial Intelligence Journal*, 172(14):1700–1729, 2008.

J. W. Weibull. *Evolutionary Game Theory*. The MIT Press, 1996.

M. P. Wellman. Methods for empirical game-theoretic analysis. In *Proceedings of the 21st National Conference on Artificial Intelligence*, pages 1552–1555, 2006.

B. Widrow and M. E. Hoff. Adaptive switching circuits. *IRE Wescon Convention Record*, 4: 96–104, 1960.

P. R. Wurman, M. P. Wellman, and W. E. Walsh. Specifying rules for electronic auctions. *AI Magazine*, 23(3):15–23, 2002.

D. Yuen, A. Byde, and N. R. Jennings. Heuristic bidding strategies for multiple heterogeneous auctions. In *Proceedings of the 17th European Conference on Artificial Intelligence*, pages 300–304, 2006.