University of Bath

UNIVERSITY OF
BATH

PHD

Evolutionary mechanism design using agent-based models

Li, Xinyang

Award date:
2012

Awarding institution:
University of Bath

Link to publication

Evolutionary mechanism design using agent-based models

Xinyang Li

A thesis submitted for the degree of Doctor of Philosophy

University of Bath

School of Management

May 2012

Contents

List of Figures

List of Tables

Acknowledgements

I would like to take this opportunity to thank all the people who made this work possible.

First and foremost I would like to express my sincere gratitude to my supervisor Dr. Andreas Krause for the continuous support, encouragement, invaluable suggestions, and great efforts to teach, explain and help. His guidance helped me in all the time of research and writing of this thesis. Without his supervision, this work would not have been a reality.

I would also like to thank Dr Richard Fairchild and Professor Edward Tsang for agreeing to be on my thesis committee and reviewing my thesis. A special thanks to Dr Richard Fairchild, who provided a number of constructive comments at the transfer viva, and which have contributed to the final thesis.

I am heartily thankful to the Scholarship Committee of the University of Bath, who awarded me with the University Research Studentship and the ORS Award. I am very grateful for the financial support which enable me to continue my research.

I would like to extend my gratitude to all my colleagues and the staff in School of Management for the help given. A very special thanks to Ronghui Miao for his invaluable advice, encouragement, and support. I would also like to thank my friend Guangyu Geng for useful advices on Matlab programming.

Lastly, and most importantly, I would like to thank my parents, Yanping Li and Hualin Wang, for their unconditional support and encouragement. In particular, the patience and understanding shown by them during the past few years are greatly appreciated. To them I dedicate this thesis.

Abstract

This research complements and combines market microstructure theory and mechanism design to optimize the market structure of financial markets systematically. We develop an agent-based model featuring near-zero-intelligence traders operating in a call market with a wide range of trading rules governing the determination of prices, which orders are executed as well as a range of parameters regarding market intervention by market makers and the presence of informed traders. The market structure which generates the best market performance is determined by applying the search technique Population-based Incremental Learning, guided by a number of performance measures, including maximizing trading volume or price, minimizing bid-ask spread or return volatility.

We investigate the credibility of our model by observing the trading behavior of near-zero-intelligence traders with stylized facts in real markets. Based on computer simulations, we confirm that the model is capable to reproduce some of the most important stylized facts found in financial markets. Thereafter, we investigate the best found market structure using both single-objective optimization and multi-objective optimization techniques. Our results suggest that the best-found combination of trading rules used to enhance trading volume may not be applied to achieve other objectives, such as reducing bid-ask spread. The results of single-objective optimization experiments show that significantly large tick sizes appear in the best market structures in most cases, except for the case of maximizing trading volume. The tick size is always correlated with the selection of multi-price rules. Though there is no particular combination of priority rule and multi-price rule achieving the best market performance, the time priority rule and

the closest multi-price rule are the most frequently obtained rules. The level of market transparency and the extend of market maker intervention show ambiguous results as their representative parameter values change in a wide range. We also find that the results of multi-objective optimization experiments are much similar to those obtained in the single-objective optimization experiments, except for the market transparency represented by the fraction of informed trader, which shows a clear trend in the multi-objective optimization. Using the results obtained from this research we can derive recommendations for exchanges and regulators on establishing the optimal market structure; for securities issuers on choosing the best exchange for their listing; and for investors on choosing the most suitable exchange for trading.

Related Publications

Below are the published papers related to this research. The first paper corresponds to section 4.3, the second paper relates to appendix A, the third paper connects with section 5.2.1, and finally, the last paper relates to section 6.2.1.

1. Xinyang Li and Andreas Krause, A Comparison of Market Structure with Near-Zero-Intelligence Traders, Intelligent Data Engineering and Automated Learning 2009, Lecture Notes in Computer Science Volume 5788/2009, pages 703-710, Springer-Verlag, 2009.

2. Xinyang Li and Andreas Krause, An Evolutionary Multi-objective Optimization of Market Structures Using PBIL, Intelligent Data Engineering and Automated Learning 2010, Lecture Notes in Computer Science, Volume 6283/2010, pages 78-85, Springer-Verlag, 2010.

3. Xinyang Li and Andreas Krause, Determining the optimal market structure using near-zero-intelligence traders, Journal of Economic Interaction and Coordination, Volume 5, No. 2, pages 155-167.

4. Xinyang Li and Andreas Krause, An Evolutionary Multi-objective Optimization of Trading Rules in Call Markets, Intelligent Systems in Accounting, Finance and Management, Volume 18, Issue 1, pages 1-14, 2011.

1. Introduction

1.1   Research background

Financial markets, in particular stock markets, have grown significantly in both developed and developing countries over the last decades. Rising globalization has aided in their growth. A large number of companies seek additional foreign listings rather than merely choosing the stock exchange in their country of incorporation, see Pagano et al. (2001), Pagano et al. (2001), Karolyi (2006). Such foreign listings were often hindered by different regulatory regimes, most notably accounting standards and disclosure rules. The requirement to present all documents in the local language will in many cases also have put significant additional costs on such a listing, easily outweighing any benefits. The convergence of accounting standards, disclosure rules and acceptance of documents in English have reduced the barriers of overseas listings considerably. Additional listings abroad have become much more common, further accelerated by improved communications technologies that enable companies to disclose any information simultaneously at any exchange.

Similarly investors are no longer restricted to the exchange in the country in which they are located. Liberalized investment restrictions for institutional investors, especially pension funds, and improved communication technologies in combination with reduced brokerage costs have made trading on exchanges abroad much more attractive. Consequently the competition has further increased through the creation of Electronic Communications Networks (ECNs) which have been set up by large investment banks to trade

among themselves at lower costs than an exchange could offer, see Benhamou and Serval (2000), Noia (2001), Ramos (2003) and Barclay et al. (2003).

Given the development and changes, exchanges have begun to compete for both listings by companies and for trading by investors. They reacted to this increased competition by investing into technology and allowing more and more electronic trading as well as electronic access to the trading floor, see Economides (1995), Barth III et al. (2002), Weber (2006). As reviewed by Economides (1995), since its first application in stock exchange in 1977, electronic trading system has been widely used in Tokyo (1982), Paris (1986), Australia (1990), Germany (1991), Israel (1991), Mexico (1993), Switzerland (1995), and elsewhere around the globe. While these changes were able to make trading faster and more reliable, this was not enough to attract sufficient trading volume. Market microstructure theory [1], which concerns the detailed process and outcomes of how exchange occurs in markets, suggests that the trading rules applied by a market affect determinants of prices, quotes, transaction costs, volume, and trading bahavior, see O'Hara (1995) and Madhavan (2000) for an overview. Changes to the market structure such as reduced tick sizes, and changes to the handling of limit orders aimed to reduced the costs of trading for investors. Thus, the market structure of an exchange becomes increasingly important in a more competitive environment.

When conducting any adjustments to their market structures, exchanges should be guided by theoretical as well as empirical research. However, the ability of market microstructure theory to provide insights into the optimal market structure is limited. First of all, the sizeable literature on auction markets can not easily be transferred to the call auction or double auction mechanisms used in financial markets. Inspired by the auction of 3G telecommunications licences the literature usually assumes the existence of a single seller and a small number of potential buyers. The buyers are furthermore assumed not only to compete in the auctions but use the acquired good from

---

[1] See section 2.1 for a detailed review.

the auction in competition outside the auction. Such a scenario is very different from the financial stock markets and is thus of limited use. In addition, as each market structure requires a slightly different model to allow it being analyzed analytically, comparisons between two or more market structures are limited. Given the large number of possible market structures, any optimization procedure involving different models would be unattainable. Moreover, most of the models in the current microstructure theory investigate the impact of a single or a small number of related trading rules on specific aspects of market performance, ignoring hybrid systems which are much more difficult to model. Furthermore, when using different models for different market structures it is even more difficult to distinguish between the effect attributable to differences in the market structures and to different behaviors of traders. These behavioral assumptions in those models make it even more difficult to assess the impact the changed trading rules have on the outcome, relative to behavioral influences. Therefore, if we seek to explore the optimal market structure, it is essential to apply a single generalized framework for all market structures considered with a wide range of trading rules applied and minimum influence from the traders' behavior.

1.2   Research aim

Although lots of stock exchanges are striving to improve their market structure, not many studies investigate the mechanism design with market microstructure theory in an integrated framework. The goal of this research is to complement and combine market microstructure theory and mechanism design in a first attempt to optimize the market structure of financial markets systematically across a wide range of trading rules. The optimal market structure is determined not only for a variety of performance measures but also a range of trader and market characteristics. We aim to evaluate the robustness of the optimal market structure to changes in the environment as well as the relevant performance measure. From these results valuable

insights can be gained on the optimal market structure for financial markets whose traders have given characteristics.

Thus far only very limited attention has been paid to the optimization of financial markets in the agent-based literature. In a series of papers Phelps et al. (2002), Phelps et al. (2003), Phelps et al. (2005), Niu et al. (2008), and Phelps et al. (2010), amongst others, use genetic programming to search for optimal trading rules in markets, but allow for trading strategies by at least partially rational traders to evolve concurrently in a co-evolutionary process. Optimal security design in a setting with adaptive agents has been investigated by Noe et al. (2003) and Noe et al. (2006). The rationality of traders and their learning in these models is, however, markedly different from the approach taken in this paper with its ZI traders; these contributions furthermore commonly investigate only the price setting behavior rather than a full set of trading rules that are available in real markets and are the focus of our investigation here. Therefore, to develop a generalized framework for determining the optimal market structure of financial markets, we conduct a research into the evolutionary design of a call auction mechanisms by investigating a large number of trading rules simultaneously. Optimal market structures for different market participants are obtained by applying the search technique Population-based Incremental Learning (PBIL). PBIL is a recently developed evolutionary optimization algorithm in which the genotype of the whole population is evolved rather than individual chromosomes. This algorithm, proposed by Baluja (1994), has been found to be simpler and to achieve better results than the standard genetic algorithm in many circumstances, see, e.g., Baluja (1995). We investigate a wide range of trading rules that jointly establish the market structure, including the tick size, degree of intervention by market markets, multi-price rules, priority rules, and market transparency, the intention is to obtain the optimal combination of these trading rules for a wide range of market characteristics.

In order to overcome the difficulties arising from conventional models in the market microstructure literature, we develop an agent-based model in which traders use a very simple trading algorithm which does not assume rational

behavior or any other optimizing rule. Such zero-intelligence (ZI) traders have been first introduced in Gode and Sunder (1993) with the explicit aim to investigate the importance of the trading rules for the outcomes of trading. The strategic behavior has been considered to be a dominant influence factor for the market dynamics in previous research. However, Gode and Sunder (1993) find that many of the major properties of double auction markets including the high allocative efficiency are primarily derived from the constraints imposed by the market mechanism, independent of traders' behavior. In Cliff (2001) such traders have also been used to determine the optimal type of auction market. The use of appropriate automatons would allow us to focus on the influence the market structure, i.e., the set of trading rules, has on the outcomes.

Market participants often value different aspects of an exchange and will thus like to use different performance measure. Given that all types of market participants will be active in the market, all their concerns have to be taken into account when determining the optimal market structure. We intend to address the problem by employing a wide range of performance measures as objective functions for the optimization, including the volatility of prices, market liquidity and trading costs, representing the different interests of various stakeholders. We investigate not only single-objective optimization but also multi-objective optimization, which allows us to determine the most appropriate sets of rules a market should consider without facing the need to balance the different interests at this stage. Using computer simulations of the trading behavior, the market performance can easily be determined for a variety of performance measures.

The results we obtain will increase our knowledge of optimal market structures in financial markets, complementing the results obtained in market microstructure theory and mechanism design. Using the results obtained from this research it is possible to derive recommendations for exchanges, regulators on establishing the optimal market structure, for securities issuers to choose the best exchange for their listing and for investors to choose the most suitable exchange for trading. While this research project will only

be able to address generic scenarios and derive conclusions on the optimal market structure from them, the developed framework will allow tailoring the circumstances to the specific needs of an exchange as well as other stakeholders and give specific recommendations to improve the market structure.

1.3 Thesis outline

Having presented the research background and the intention of the work, the remainder of the thesis is organized as follows:

Chapter 2 provides an literature review of the market microstructure theory and the market simulation techniques. In section 2.1, we will give a brief overview of the trading process in stock markets, the market forms that exist in the financial market, and various orders types and their submission rules. This is followed by a review of a set of trading rules that are applied in financial markets. We then review three common financial markets in more details, namely the dealer market, auction market, and limit order market. In section 2.2, we discuss some models applied in the financial market simulation, including the agent-based models, zero-intelligence models, and zero-intelligence-plus models.

Chapter 3 will give an overview of market mechanism design. At first We will present a review of auction theory and auction design, followed by a discussion of simulation studies in mechanism design. In section 3.4 we combine market microstructure theory and mechanism design, and discuss several market microstructural factors which need to be considered while optimizing the market structure.

In chapter 4, we study the evolutionary optimization techniques for both single-objective optimization and multi-objective optimization. For single-objective optimization, population-based incremental learning is considered, which is an evolutionary optimization algorithm that combines the concept of

competitive learning and genetic algorithms. Before discussing the evolution-
ary multi-objective optimization algorithms and their empirical comparison,
some basic concepts and terminology of multi-objective optimization are in-
troduced, followed by a review of two traditional approaches, which are the
weighted-sum method and the normal constraint method.

In chapter 5, we propose an agent-based model featuring near-zero-intelligence
traders operating in a call market with a wide range of trading rules govern-
ing the formation of prices at which orders are executed, as well as a range
of parameters regarding market intervention by market makers and the pres-
ence of informed traders. After determining the credibility of our models by
reproducing several return properties that are commonly found in financial
markets, we attempt to optimize the applied trading rules using population-
based incremental learning approach seeking to obtain the optimal market
structure with both single-objective function and multi-objective functions.

Finally, chapter 6 summarize the research, the contributions and limitations,
and some suggestions for future study.

2. Literature review

This chapter is divided into two sections. The first section reviews the market microstructure theory. Section 2.2 presents some simulated studies of financial markets using agent-based models, zero-intelligence traders, and zero-intelligence-plus traders.

2.1 Market microstructure theory

Garman (1976) first introduced the term "market microstructure", as the title of a study. In the research, Garman analyzed the effects of market structures and trading behavior on the process of price formation in the work. Market microstructure has then become a descriptive title for the study of the trading mechanisms used for financial securities. It studies trading process, and structure of markets defined by various trading rules and trading systems, e.g., which agent can trade, at what time, in what manner. As later described by (O'Hara, 1995, p 1),

> "Market microstructure is the study of the process and outcomes of exchanging assets under explicit trading rules. While much of economics abstracts from the mechanics of trading, microstructure theory focuses on how specific trading mechanisms affect the price formation process. These mechanisms may involve a specific intermediary such as a stock specialist or an order clerk (a saitori), employ a centralized location such as an exchange or a futures pit, or be simply an electronic bulletin board in which

> buyers and sellers indicate an interest in trading. Whatever the specific mechanism, however, prices emerge and buyers and sellers trade."

In the last two decades, there has been a tremendous growth in the academic literature on market microstructure. Most research explores the impact of a specific trading mechanism on price behavior, or compares the market performances under different circumstances. It studies the market behavior to explain issues, such as how trading prices are determined, what properties the prices exhibit, why they emerge, and how trading rules evolve in markets, which have important implications for mechanism design and market regulation. The microstructure of markets can directly affect the selection criteria of investors when choosing a market, such as liquidity, return and risk. To attract investors, markets should be more efficient and competitive: allowing buyers to pay less, sellers to obtain more, and providing liquidity inexpensively. In order to effectively design or optimize a market mechanism, we need in the first step to understand the market structure, which provides the framework within which the market operates. Therefore, a general market design problem has become to design a market mechanism where agents interact and trade.

This section reviews the literature on market microstructure theory. It is not the aim to provide a complete survey of all models and trading systems, but to explain the basic principles. After some preliminary definitions in the first section, we will then describe several different forms of market structures in section 2.1.2. Section 2.1.3 shows the order submission process and various order types, followed by a discussion of impacts of a set of trading rules applied in financial markets. Finally, three typical markets, including dealer market, auction market, and limit order market, are reviewed in the rest of this part.

2.1.1 Definitions

In economics, a *market* is any place allowing buyers and sellers to exchange any goods, services and information, while a *financial market* is a mechanism allowing market participants to trade financial securities, commodities and other transferable items at very low cost.

For the *market participants*, O'Hara (1995) shows that there are four main categories[1]. First of all, there are traders who place orders in the market. An *order* is an instruction from a customer to a broker to buy or sell a certain quantity of assets. Next, there are *brokers*, who do not trade for themselves, merely transmitting orders for customers. Thirdly, there are *dealers* who are ready to buy and sell for their own account. Finally, there are *market makers*, or *specialists*, who do not only take a position in the securities, but also quote purchasing prices or selling prices. As noted by O'Hara (1995), "Since the market maker generally takes a position in the security..., the market maker also has a dealer function. The extent, however, to which the market maker acts as a dealer can vary dramatically between markets."

Securities are fungible and negotiable instrument with financial value or rights with a sequence of future cash flows consisting of money or other securities. They are categorized into debt securities such as equity, bonds and debentures, banknotes, and derivatives. *Bid price* is the highest price at which a buyer is willing to pay for a particular good. It is also referred to as the "bid". On the other hand, the best price that a seller is willing to sell is the *ask price*. The difference between the bid price and ask price is called *bid-ask spread*.

2.1.2 Market forms

Before reviewing issues of trading process in different market structures, it is important to understand what market structures encompass and the major

---

[1] see (O'Hara, 1995, p 8)

types of market structure, even though most realistic financial markets are mixed market structures.

As summarized by Madhavan (2000), *market architecture*, or market structure, is determined or distinguished by the selected set of rules applied during trading in respect of (1) *market type*: with the consideration of the *degree of continuity* of the trading system, *reliance on market makers* and the *degree of automation*; (2) *price discovery*: which measures to what extent the transaction price is discovered by an independent price formation process or determined by another market; (3)*order forms*: such as market orders, limit orders, and stop orders; (4)*protocols* or trading rules: for example, choice of tick size and priority rules employed to determine the rationing of orders in case an imbalance between buy and sell order exists at the transaction price, and (5)*transparency*: the quantity and quality of trading information disseminated in the market. This taxonomy of market structures will help us distinguishing various market forms.

The *Walrasian tâtonnement* or *Walrasian auction* has been a fundamental assumption in conventional financial theory since Leon Walras introduced the general equilibrium theory in 1874. In this mechanism, each trader determines his demand or the amount he is willing to buy or sell at a stated price and submits orders to the *Walrasian auctioneer*, who aggregates traders' demand and supply and determines the clearing price. After the auctioneer suggests a potential price for trading, traders then revise their demand to the optimal level at the stated price. This process continues until there is no further revision and the aggregate demand is exactly the same as the aggregate supply. The price at which each trader submits his optimal order is called the *equilibrium price*. All the submitted orders are executed in a single transaction at the equilibrium price.

In the Walrasian framework, the process of price formation is captured by the representation of a Walrasian auctioneer, and the trading price proceeds from a tatonnement, i.e., a series of revisions. No trading exists outside the equilibrium price. This auction activity is costless as the process is assumed

to be finished in an instant of time. However, the revision of orders and the determination of clearing prices would take a long time in real world, especially with more traders. This would definitely impose a cost during trading. Moreover, it would be very difficult to find a clearing price that balances the aggregate demand and supply. Because of these limitations the Walrasian framework does not often exist in financial markets, instead, a wide variety of some other market structures have been formed, see Krause (2000) for more details of various market forms.

As stated by Krause (2000), "In all market forms orders can be submitted to the market at any time, what differs is the time and way these orders are executed." A market with similar concept of an Walrasian auctioneer is the *call market* or *batch system*. In this market, all the bid and ask orders are aggregated and transacted at once at a predetermined point of time. The price is set by the exchange based on the bid and ask orders so most orders can be executed and the market will clear or almost clear.

Rather than executing orders in predetermined time intervals, in *continuous markets*, trades can take place at any point of time. As explained by Krause (2000),"For every submitted order it is immediately checked whether there exists another order on the market, such that these orders can be executed in a bilateral trade. If no such order exists, the order is stored and executed with the next matching order arriving on the market."

A *dealer market* is a widely used market form for financial assets transactions. In this market, securities are traded through a network of dealers who provide liquid by buying and selling securities with traders. The market maker set prices at which he is willing to buy and sell the assets. In the United States, for example, bonds and foreign exchange are traded almost exclusively in dealer markets.

Another widely applied market structure is called *auction market*. It is a system in which buying and selling financial assets take place on the floor of an exchange, such as the New York Stock Exchange. The price is determined through the open and free interaction between traders.

In *limit order book market* all bids and asks submitted in the market are gathered and kept in an order book. A trade occurs when two orders cross, for example, if the limit price of an incoming ask order is less than the limit price of a bid order, the orders are immediately traded.

Amihud et al. (1990) compare the performance at the opening of a continuous auction market and a call auction market. The call auction has been found to be more effective with lower volatility, and efficient prices which is determined at the opening while information is highly asymmetric. Consistent with this result, Madhavan (1992) also observes that a call auction mechanism "can function where a continuous market fails" when the problems of information asymmetry is severe.

Economides (1995) propose to apply an electronic call auction mechanism three times a day in a market: at market open, at noon, and at market close. The call market clears a bunch of orders at predetermined points of time, which improves the liquidity. As a result, the transaction costs for public participants are reduced and the market efficiency is improved. Similar research has been done by Rhee et al. (2004) and Comerton-Forde et al. (2005), who study the impact of call mechanism at the market open and close on the Singapore Exchange. Rhee et al. (2004) argue that price discovery process is improved, and market manipulation is decreased at the market close with the introduction of call market. Comerton-Forde et al. (2005) find that call auctions significantly increase the trade volume, and simultaneously reduce the incidence of closing price manipulation. Pagano and Schwartz (2003) report that the use of call auction mechanism at the market close on the Paris Bourse improves price discovery and reduces execution costs significantly.

2.1.3 Order submission and order type

An order is an instruction from traders to brokers to buy or sell in a market. In markets, orders are not traded directly between investors. They are first

submitted to a broker, who then transmits the orders to the market. When the order is executed, the broker informs the investor about the fulfillment and the transaction price of the order and settles the accounts of the customer. In a dynamic market, the price paid or received may not always be the last quoted price before the order was entered. Once the market order is placed, the investor has no control over the price at which the transaction is executed.

The order submission process might be different in various market structures. For example, in batch systems the broker transmits the order to a *match maker*, who keeps the order book and matches the orders, instead of transmitting them to an exchange. The order size is not fixed in markets with the exception in board trading. A market order for a large number of shares is normally divided into several small orders by the broker and executed over a longer period, resulting in different prices for parts of the order.

There are several forms of orders. The simplest type is *market order*, which is an order to buy or sell immediately at the current market price. The advantage of a market order is that it is always guaranteed to be executed as long as there are investors willing to trade. Another frequently observed order type is *limit order*. It is an order to buy or sell a security at no more or no less than a specific price. A buy limit order is executed at the limit price or lower; while a sell limit order is executed at the limit price or higher. It avoids buying or selling securities at a too high or low price. Comparatively, it gives investors some control over the price at which the transaction is taken. However, the transaction of limit orders is not guaranteed as market orders. The price may surpass the limit before the order can be executed. Furthermore, in some markets transactions between market orders have higher priority than transactions between a market order and a limit order. Because of the complexity, it may take quite a long time to be settled. Apart from the market order and limit order, there are some other rarely applied types of orders, such as the stop loss order, market-if-touched order, one cancels other order, and tick sensitive order.

2.1.4   Trading rules

Market microstructure theory as used in conventional finance suggests that the trading rules applied by a market affect the market structure and market performance, in particular the prices at which trades occur. This influence on prices should also be visible in the statistical properties of returns such as their distribution and autocorrelations. Intensive debate on how trading systems should be designed and the impact of different design features has been sparked. This section provides an overview of some important design features, including tick size, priority rules, multi-price rules, the degree of transparency and the intervention of a market maker.

2.1.4.1   Tick size

The *tick size* is the minimum differences between prices at which orders can be submitted, usually determined by the exchange where the order placed. On the New York Stock Exchange (NYSE), the tick size was 1/8 dollar prior June 1997. Afterwards it was reduced to 1/16 in 1997 and further dropped to one cent in 2001 in US equities markets. In dealer markets, a standard tick size does not always exist. NASDAQ, as an example, does not have a market-wide standard tick size. In futures markets, each contract has a prescribed tick size that relies on the future contract value and its variability. For instance, the S&P 500 future contract has a tick size of 25 dollar per contract or 0.10 index points.

The tick size has several impacts during trading. First of all, as it represents the cost of getting inside other competitors' quote, the tick size affects the motivation of submitting limit orders.[2] For example, if the tick size is 1/8 dollar or 12.5 cents, and the standing bid is 10 dollar, a trader must place an order with the price 10.125 to compete and move in front of the standing bid. If the tick size is one cent, the trader must place an order with the

---

[2] First noted by Harris (1991).

price of 10.01 to move in front of the standing bid. Since when the tick size decreases, it becomes easier to be ahead of another limit order, fewer orders are needed, which could have adverse impacts on liquidity. Therefore, the decreased tick size reduces the depth of the market, measuring the size of the best limit order, see Ricker (1999) and Huang and Stoll (1999). In this case, traders placing small orders benefit from the small tick size while large orders are traded at less favorable prices.

Some other papers, e.g., Harris (1994), Allaudeen and Eric (1998), and Chan and yang Hwang (1998) study the impact of tick size on trading volume. Harris (1994) find that the trading volume would be enhanced (by 34%) with a decreased tick size (to one-sixteenth for NYSE stocks). Allaudeen and Eric (1998) suggests that trading volume is more likely to increase with a reduced tick size if the stock is actively traded. Besides, research, such as Harris (1990a), Ronen and Weaver (1998), and Bessembinder (2000) confirms that observed volatility decreases with a reduction in tick size.

In addition, the tick size has an impact on the spreads. If the tick size is 12.5 cents, the minimum spread between quotes is 12.5 cents. This spread may surpass the equilibrium spread and cost more to pay. Harris (1994) and Ball and Chordia (2001) investigate true spreads and price behavior with artificial price increments imposed by the minimum price variation. As reported by empirical investigation of NASD (1997), Ricker (1999), Goldstein and Kavajecz (2000), and Jones and Lipson (2001), the spread declined dramatically by about 25% with the reduction of tick size from 1/8 to 1/16 dollar.

Tick size is an important variable of choice for mechanism designers. It explains why share price levels differ substantially across markets or countries. It also affects the trading costs since it sets a lower limit on the minimum size of the bid-ask spread. Angel (1997) investigates optimal relative tick sizes and demonstrates that they are a function of firm size, visibility of the firm, and its idiosyncratic risk. He argues that firms may split their stock to move their share prices into the range where the mandated tick size is optimal relative to the share price, which implies that firms would have an

opportunity to choose their own tick size. In addition, he mentions that "A wider tick size enhances liquidity by reducing bargaining and processing costs and by providing more incentives for limit orders and market makers to provide liquidity ...", however, it "also increases the minimum quoted bid-ask spread."

Tick size in many stock markets has been lowered[3]. A large number of empirical papers examine the influence of reduced tick sizes. Ahn et al. (1996) and Bacidore (1997) study the reduction of tick size on the Toronto Stock Exchange. They argue that a smaller tick size leads to narrow spreads together with a decrease in quotation size. Consistent with this result, Lau and McInish (1995) observe that a reduction of tick size from $0.5 to $0.1 for shares at $25 or more on the Singapore Exchange lowers the bid-ask spread. Ke et al. (2004) examine the adverse scenario where tick size rises from $0.1 to $0.5 for shares above $50 on Taiwan Stock Exchange. The result shows that both spreads and volatility increase. Nevertheless, it is also found that if tick sizes are too narrow market quality may be harmed, see Harris (1994) and Demarchi and Foucault (2000). Harris (1996) argue that with a large tick size limit order traders are protected against the free option problem and traders are induced to post more orders with larger sizes. Similarly, Cordella and Foucault (1999) find that more rapidly competitive orders are placed with a larger tick size. Huang and Stoll (2001) indicate that the tick size is interconnected with some other features like market depth and spreads and together reveal the underling market design. They also suggest that in order to provide more liquidity in auction markets, it is necessary to have a minimum tick size, otherwise a limit order could easily step ahead of the dealer's quote or another limit order. For other related papers, see Darley et al. (2000), Chiarella and Iori (2002), and Yeh (2003).

---

[3] For example, the AMEX in 1992, the Singapore Exchange in 1994, the Toronto Stock Exchange in 1995, and NYSE in 1997

2.1.4.2 Priority rules

Another important design feature of trading systems is the enforcement of priority rules, which determine the execution sequence of orders in the market. Assume a trader wants to purchase 100 shares of stock and many sellers are willing to sell in the market. The problem is which seller should get to trade with the buy orders. In this case markets use a variety of priority rules to match submitted bids and offers, see Domowitz (1993) and Weaver (2006) for an overview of the different priority rules found in several markets.

On all major world markets, such as London Stock Exchange (LSE), the NYSE and Euronext, price priority is enforced. With the *price priority*, buyers who are willing to pay the highest price will be the first to buy and sellers willing to sell at the lowest price will be the first to sell. However, it is often found that more than one buyer is willing to buy at the same price. We then distinguish between the unexecuted orders with *secondary priority rules*.

The secondary trading priority rule [4] defines the sequence to be followed for orders with the same price. The *time priority* rule, which represents first-come, first-served, is the most commonly used rule. It ensures that the order which is submitted earlier at a given price is executed first. Many exchange markets, including Paris Bourse, Tokyo Stock Exchange, and the American Stock exchange, apply some variation of this rule.

Another frequently used rule is *size priority*, which grants priority to orders with the largest size. It has the advantage of promoting traders to place larger orders. Alternatively, the order which matches the incoming order in size could to executed with priority. A variation of this rule is applied on the New York Stock Exchange and Toronto Stock Exchange (prior to 1996).

Pro rata priority is also a common practice on many financial markets such as the Stock Exchange of Hong Kong and the old Toronto Stock Exchange

---

[4] See (Schwartz, 1988, p 18)

with floor-based trading system. It allocates equal shares to each order on the larger side if there is an imbalance in the orders at a given price. This rule could also be used in batch systems.

Another rationing rule is *random priority* which randomly selecting orders at a certain price. Each trader willing to trade at that price has the same probability of filling the order. This is similar to the "open outcry" method in floor-based system such as the Chicago Board of Trade.

There are some other rules that are not applied very often. For example, *class priority* rule selects orders based on the classes of traders, and *exposure priority* rule gives priority to orders that are revealed to other market makers.

It has been found that the various priority rules applied have different impact on trading costs, price and liquidity. Harris (1990b) and Angel (1997) indicate that the choice of secondary priority rules is vital due to the impact on market efficiency and traders' motivation to provide liquidity. Harris (1990b) argues that secondary time priority encourages price competition among traders while secondary size priority discourages traders from providing liquidity. He argues that price should be the dominant priority, followed by orders displaying broker identity and size, then orders based on their arrival time. In line with Harris (1990b), Demarchi and Foucault (2000) suggest to adopt price then time priority to support price competition, since traders who submit late orders must offer better prices than existing quotes to enhance the execution probability. By using a simulated queuing model, Cohen et al. (1985) find that time priority leads to more price competition, and bid-ask spreads in markets that employ time priority are smaller than those that do not enforce time priority. With a theoretical model of dealer competition, Cordella and Foucault (1996) find that spreads is narrower in markets in which price/time priority rule is enforced than those under a random selection rule.

Angel and Weaver (1998) and Panchapagesan (1997) study the differences of market quality and trader behavior between markets under time priority and under pro-rata priority rule. Both papers provide evidence that limit order

traders compete more on prices if time priority rule is employed. In addition, if priority is based on size instead of time, limit order traders would indulge in gaming strategies such as submitting larger orders in an attempt to increase the probability of execution. Panchapagesan (1997) further notices that lack of price competition under sharing priority rule leads to wider spreads than under time priority rule.

2.1.4.3 Multi-price rules

In some markets, the transaction price is determined where the demand and supply curves intersect, i.e., the price at the maximal trading volume is chosen as the transaction price. If there are multiple prices at which the trading volume shows the same maximal value, multi-price rules will be employed to determine which of the prices will be chosen. It could be either the price closest to the previous price, the price furthest from the previous price, the highest price, the lowest price, the price with the highest order volume, the price with minimum order imbalance (the absolute difference between the aggregate order size of buy and sell orders at the transaction price), the price with maximum order imbalance or a randomly selected price. Each market has its own multi-price rule. However, little attention has been paid to the influence of different multi-price rules on the trading performance.

In the real world trading markets, for example, Frankfurter Wertpapierbörse (FWB) enforces the multi-price rule of choosing the price which allows for the greatest transaction volume (principle of maximum execution)[5]. Similarly, Shanghai Stock Exchange (SSE) uses the price that minimize the unexecuted volume as the execution price. If there are multiple prices that minimizes the unexecuted volume, the mean price is calculated and applied as the execution price.

---

[5] See $http://deutsche-boerse.com/dbag/dispatch/en/kir/gdb\_navigation/info\_center$

2.1.4.4 Transparency

A large number of informational issues regarding market microstructure concern information and disclosure. *Market transparency*, as defined by O'Hara (1995), is "the ability of market participants to observe the information in the trading process." In this context, information refers to knowledge about the prices, the size and direction of orders, the order form, and the identities of market participants. Some developed automated trading markets display the entire contents of the limit order book, which are totally transparent. Other markets, on the other hand, permit lower degree of transparency. They display only quotations with the highest bid or lowest offer. In the real-world markets, for example, the ECNs disclose information of the whole book, while the NYSE only discloses information of orders at the top of the book, i.e., the best bid and ask.

Transparency in the markets can be divided into two discrete components: pre-trade and post-trade transparency. *Pre-trade transparency* refers to the information available to market participants ahead of trading, indicating the price and size of prospective trading interest, such as bid and ask quotations, depths, and resting limits orders, both at the best quotations and those away from them, as well as other trade related information such as the order imbalances. *Post-trade transparency* refers to the public dissemination of information on past transactions, containing the execution time, trade price, volume and probably the identifications of traders.

In a transparent market, traders are able to have access to information on the order book. This could reduce the magnitude of adverse-selection problems, hence transparency is expected to increase profit for traders. A large academic literature has studied the importance and the optimal amount of information disclosure before trading. Biais (1993) argues that quotation transparency enhances market efficiency and liquidity. Lyons (1994) shows that little disclosure in foreign exchange markets causes a "hot potato" effect, which generate additional volatility. Pagano and Roell (1996) provide a contrast of several different types of trading systems, and indicate that

uninformed investors prefer transparent markets because of the higher liquidity and lower trading costs. Flood et al. (1999) discover that the pre-trade transparency can narrow the spreads. Boehmer et al. (2005) examine the introduction of NYSE OpenBook, which make the Specialist's limit order book public to all traders. Similar findings of increased liquidity and returns, and reduced execution costs have been obtained in a more transparent market. In addition, increasing transparency leads to more aggressive limit order submission and enhancement of market depth. And, Hendershott and Jones (2005) notice that a decrease in transparency increases trading costs and exacerbates price discovery. As demonstrated by Goldstein et al. (2007), with increased transparency, "Measures of trading activity, such as daily trading volume and number of transactions per day, show no relative increase, indicating that increased transparency does not lead to greater trading interest". In addition, Aitken et al. (2001) confirm that "the enhancement in pre-trade transparency, through tightening the undisclosed order regulation in October 1994, resulted in a significant decline in trading volume."

However, one can argue that transparency can make it difficult to supply liquidity to large traders, who may be reluctant to submit limit orders, since the disclosure may convey information which makes the price moves against the traders' position. Madhavan (1996) argues that greater transparency raises price volatility, and reduces liquidity, which harms the market quality. However, in a sufficiently large and liquid market, transparency has been found to decrease volatility, increase market liquidity, and lead to more stable prices, since traders in a very liquid market may not change their strategies significantly based on order book disclosure. Porter and Weaver (1998b) study the impact of a higher degree of transparency on the Toronto Stock Exchange with the transmission of the best quotes and associated depth as well as quotes for up to four levels away from the inside market in both directions. They have found that both effective spreads and the percentage bid and ask spread increase with the introduction of the system, i.e., pre-trade transparency decreases liquidity. Furthermore, Madhavan et al. (2005) report that increased transparency by providing order book disclosure on

the Toronto Stock Exchange in 1990 induces higher trading costs and lower liquidity.

Other studies such as Chowdhry and Nanda (1991), Baruch (1997), Bloomfield and OHara (1999) and Flood et al. (1999) reach mixed conclusions regarding the effects of transparency. For instance, Bloomfield and OHara (1999) demonstrate that even though greater transparency leads to an increase in the informational efficiency of trade prices, it enlarges spreads at the same time.

Another issue of pre-trade transparency on whether to display brokers' identity is still on debate. The degree of anonymity is critical because it can influence the level of information dissemination to which informed traders benefit from their informational advantages, and thus affect the adverse selection costs. Moreover, the level of anonymity have effects on both price formation and liquidity in trading systems.

Some theoretical papers tend to argue that more informed trading will be induced in anonymous trading mechanisms, see Roell (1990), Fishman and Longstaff (1992), and Forster and George (1992). Harris (1997) notices that anonymity facilitates the management of order exposure risk by concealing traders' intentions. In particular, large anonymous orders are often considered to be placed by informed traders who prefer not to be identified (Harris (1996)). This is empirically supported by Grammig et al. (2001), Heidle and Huang (2002) and Barclay et al. (2003), who study investors' decision of whether to select anonymous or non-anonymous trading systems. In contrast, some academic literature on market transparency shows that uninformed traders can benefit from a highly transparent market, yet at a cost to informed traders. Forster and George (1992) show that higher pre-trade transparency lowers transaction costs to investors who can identify liquidity-motivated orders. Theissen (2003) suggests that broker identification should be public, as adverse selection risk is higher in an anonymous market. That is, uninformed traders are likely to suffer from adverse selection risk in an anonymous market due to the lack of information of the identity of the party

on the other side of the trade, which in turn may deter uninformed liquidity. On the other hand, Foucault et al. (2003) examine the influence of a shift to pre-trade anonymity on Euronext Paris. They find that the quoted spreads and price informativeness significantly decline after the transformation. Similarly, Comerton-Forde and Tang (2009) investigate the impacts of the removal of broker identifiers from the central limit order book of the Australian Stock Exchange. They indicate that anonymity leads to smaller spreads, higher order book depth, and lower order aggressiveness. Larger and more actives shares tend to benefit more than small shares from anonymity. The similar improvements in liquidity, quoted depth, and spreads after the removal of broker mnemonics have also been found on the Sydney Futures Exchange by Frino et al. (2008).

Apart from these, several papers, e.g., Naik et al. (1994), Gemmill (1996), Board and Sutcliffe (1995), and Saporta et al. (1999) investigate the impact of post-trade transparency on stock exchanges. Some resent papers, see Bessembinder et al. (2006), Harris and Piwowar (2006) and Edwards et al. (2007) explore the effect of post-trade transparency on bond markets. Gemmill (1996) studies the impact of delayed reporting of block trades by analyzing three different publication regimes on the London Stock Exchange. Surprisingly, he finds that delayed reporting has little effect on block spreads, speed of adjustment, volatility and even liquidity. Differing from the theoretical concept, market makers do not benefit from the delayed reporting which they request vehemently for. Waisburd (2003) studies the effects of post-trade anonymity in Paris Bourse. The result shows that liquidity in the post-trade anonymous regime is much lower than the other one.

So far, the impact of market transparency is still rather vague. As discussed above, the effects of pre-trade and post-trade transparency are different under various conditions. Thus, the problem of deciding on the optimal level of transparency is difficult to be solved considering the great differences in market mechanisms (see O'Hara (2007)).

2.1.4.5 Intervention of market maker

Except for the three design features above, a market maker would also intervene or influence the prices such that he is prepared to trade a fraction of the order imbalance at his own transaction price by submitting an offsetting order. Market makers, who are willing to buy or sell shares on demand, are essential in financial markets. They are not active to search for a matched order, but wait for orders that need to be traded. They are accountable to supply liquidity and support continuous trading by overcoming the asynchronous timing of posted orders.

Stoll (1978), Amihud and Mendelson (1980), Ho and Stoll (1981), and Ho and Stoll (1983) among others first analyze the pricing behavior of market makers. They find a linear relationship between the price set by a market maker and his/her current inventory position in pricing models of market-making, i.e., the price quoted by a market maker is determined by his/her inventory level. The general idea is that during the trade, the inventory positions diverge and force the market maker to modify the price level. As denoted by Biais et al. (2005), "Market makers with very long positions are reluctant to add additional inventory and relatively inclined towards selling. Consequently, their ask and bid prices will be relatively low. Similarly, market makers with very short inventory positions will tend to post relatively higher quotes and will tend to buy. Thus, market makers' inventories will exhibit mean reversion. Because of the central role of inventory considerations in this analysis, it is often referred to as the inventory paradigm." Therefore, increases in inventory result in drop of prices, while decreases in inventory lead to increase of prices.

A number of papers applied agent-based models[6] to study the behavior of market makers. Gu (1995) adopts the model used in Day and Huang (1990) to analyze the market maker's price adjustment rules and the profitability

---

[6] Agent-based modeling is a computational modeling for simulating actions and interactions of auonomous decision-making entities called agents. See section 2.2.2 for detailed explanation.

under different conditions. Gu argues that "to make a viable market the specialist has to chum the market, and to behave optimally he may have to sacrifice his affirmative obligation to maintain 'a fair and orderly market'". Westerhoff (2003) investigates the influence of inventory position, and argues that market maker adjust the price depending on the current inventory position and excess demand or supply. Madhavan and Panchapagesan (2000) discuss that a call auction with a dealer is more efficient than an automated call auction without a dealer. Rust and Hall (2003) find that a market trader's expected profit will be increased with participation of a market maker.

Some empirical papers discuss the role of market makers on a trading floor. For example, Madhavan and Smidt (1993), Madhavan and Sofianos (1998), Kavejecz (1999), and Madhavan and Panchapagesan (2000) explore the critical role of NYSE specialist. Academic literature on mechanism design also discuss the benefits and importance of market makers. For instance, Brusco and Jackson (1999) study the optimal design of trading rules in a two-period market. They conclude that market makers, who can economize on fixed trading costs across periods, play central roles in the optimal mechanisms. To obtain Pareto efficiency, multiple market makers are required, and the number of which is depending on factors, such as the level of risk aversion, the asymmetric information, and the discount factor.

2.1.5 Dealer market

As reviewed in the section of alternative market mechanisms, a market maker or dealer sets and electronically posts prices at which he is willing to buy or sell the security, and wait for an offsetting order to arrive and trade immediately with an investor demanding it. However, there are risks of not knowing when and at which price the market maker can clear his position. To cover the costs imposed by the risks, the market maker quotes different bid and ask prices for the security. As investors do not bear the risks, they are willing to pay for the immediacy service by trading at a less favorable price as long

as the costs from waiting an offsetting order at an unknown price equals the costs imposed by the risk taken by the market maker.[7] Therefore, a market maker is compensated by charging a fee through the difference between the ask and bid price, i.e., the spread for the services provided for enabling investors to trade immediately by forming the counterpart. The price-setting behavior among multiple competing market makers is one of the prominent features of dealer markets. The interaction of these market makers' quotes provide liquidity to investors and price formation to markets.

Due to the uncertainty that a market maker faces, concerning the future security value and order flow, the bid and ask prices setting becomes a complicated decision problem. This problem is analyzed in a multi-period framework with the knowledge of a market maker's risk preferences to determine how each uncertainty of transaction and inventory value affect the trading prices. As noted by O'Hara and Oldfield (1986), the market maker's bid and ask quotes spread is decomposed into three parts, which are a portion for limit orders that have already submitted for execution at a specific price; a risk-neutral adjustment for the submitted market orders for execution at the current price; and a risk adjustment for inventory value and order flow uncertainty.

Similarly, Stoll (1978) demonstrate that the cost of services provided by market makers is the sum of three elementary costs: order-handling costs, i.e., opportunity cost of holding an order and the price uncertainty; inventory costs, i.e., costs of recording, arranging and settling a transaction; and information costs, i.e., costs incurred if investors trade with superior information.

2.1.5.1 Order-handling cost

To simplify the analysis of Stoll (1978), Biais et al. (2005) view it through a simple synthetic model, in line with the work of Stoll (1978). Suppose there are N liquidity suppliers in the market for a risky security. Let χ be

---

[7] Waiting cost is first introduced by (Demsetz, 1968, p 37)

the expectation of the fundamental value (ν), U_i be the utility function, I_i be the information set, ς_i be the cash endowment, and γ_i be the risky asset endowment of liquidity provider i.

Biais et al. (2005) first study an instance in which market order M is submitted and then equilibrium achieved in a uniform-price auction. In this auction, liquidity trader i optimize his limit order schedule by determining the size he buys or sells $s_i(\rho)$ at each possible price ρ:

$$(2.1) \qquad \forall \rho: \ Max_{s_i(\rho)} EU_i(\varsigma_i + \gamma_i \nu + (\nu - \rho)s_i(\rho)|I_i).$$

Based on the market-clearing principle, the equilibrium price is then established as:

$$(2.2) \qquad M + \sum_{i=1,....,N} s_i(\rho) = 0.$$

In an alternative case, limit orders are placed first, and then matched with a market order. This case is considered to study the order-handling costs. Assume there are N risk-neutral market makers with the same cost $(c/2)s^2$ to make a deal of s shares, which reveals order-handling costs. Assume an uninformed trader has posted a market order to buy M shares. With the uniform-price auction model 2.1and 2.2, each market maker would sell M/N shares at the following ask price:

$$(2.3) \qquad \alpha = \chi + (\frac{c}{N})M,$$

which measure the marginal cost. Likewise, if the liquidity supplier intends to sell, the bid price is then set as:

$$(2.4) \qquad \beta = \chi - (\frac{c}{N})M.$$

The bid-ask spread is, therefore, $2(\frac{c}{N})M$. Biais et al. (2005) also mention that there would be a negative serial autocorrelation in returns with two assumptions that market orders are independent and identically distributed, and the fundamental value follows a random walk. This is caused by the bouncing of execution prices between the bid and ask quotes.

2.1.5.2 Inventory-based models

Since market makers make quotes at which they are willing to trade, they have direct influences on transaction prices. Most of the literature investigates the behavior of market makers in two major types of models, which are inventory-based models and information-based models. Inventory-based models assume that a market maker is risk-averse. All investors in the models have the same information and an agreement on the fundamental value implied by this information. On the other hand, information-based models assume that a market maker is risk-neutral. Investors in these models are divided into two groups, namely informed and uninformed investors. In both models traders are motivated by exploiting high liquidity. In addition, traders in information-based models are also motivated to exploit informational advantages. We will fist review inventory-based models, followed by the information-based models in the next part.

As first studied by Stoll (1978), the market makers are assumed to be risk averse in inventory-based models. Biais et al. (2005) review the study by simply focusing on CARA utility function and random variables from normal distribution. They suppose that the market makers' constant absolute risk aversion index is δ, P is the average inventory position of the market makers $(P = \sum_{i=1} NP_i/N)$, and σ^2 is the variance of the final cash flow of the security. Again if the liquidity supplier places a market order to purchase M shares, the competitive market makers would sell the shares at the following ask price set as the marginal valuation by using the uniform-price auction

model (2.1 and 2.2):

(2.5)
$$\alpha = [\chi - \delta\sigma^2 P] + (\frac{c + \delta\sigma^2}{N})M.$$

Similarly, the bid price is set as:

(2.6)
$$\beta = [\chi - \delta\sigma^2 P] - (\frac{c + \delta\sigma^2}{N})M.$$

As indicated by Biais et al. (2005), the midpoint of the bid-ask spread is measured by subtracting a risk premium received by the market makers to compensate for taking the risk of holding the initial inventory ($\delta\sigma^2 P$) from the fundamental value of the security (χ). The authors further explain that if a market maker takes a considerable long position, he will tend to reduce inventory instead of purchasing additional shares. As a result, the bid and ask prices will be comparatively low. Moreover, a market maker with a considerable short position will place buy orders with higher prices to build up the inventory level. Therefore, inventory level of market makers is more likely to exhibit mean reversion.

Amihud and Mendelson (1982) investigate the behavior of asset prices in a dealership market with an inventory-based models. In this model, trading prices are taking from the execution of randomly submitted orders at the market maker's buy or sell prices, which are quoted to move the level of inventory toward a preferred position. According to the research, bid and ask prices should be a decreasing function of an market maker's inventory position. They illustrate it with an example. If a market maker is at a short position, he would induce public sell orders and discourage buy orders to replenish his inventory by way of increasing bid price to increase the supply function,and by increase the ask to reducing the demand function . It is also found that the bid and ask prices quoted at the "preferred" inventory position [8] is the closest to the equilibrium price at which aggregated supply matches aggregated demand. In addition, the bid and ask spread is narrowed

---

[8] See Logue (1975)

to the least value at that position. The spread becomes increasing while the inventory position of market maker diverges from the preferred level. Overall, the quoted prices are found to be stabilized and changed moderately; and the market is efficient since investors cannot obtain any profit without superior information.

Roger and Eeckhoudt (1999) analyze the properties of bid and ask prices quoted by a monopolistic market maker without parametric assumptions about the utility function of the market maker or the probability distribution of the risky asset returns. Results show that the bid or ask price exhibits a concave or convex function of the quantity traded relatively. They prove that both prices can be higher or lower than the fundamental value of the security. Although there is not a clear shape for the bid-ask spread, the spread narrows as the inventory position is high or the market maker has decreasing absolute risk aversion, except for the situation while several market makers are competing.

In line with the inventory control literature, Manaster and Mann (1996) find that market makers do manage the inventory levels. They investigate a unique data-set describing locals trading on the Chicago Mercantile Exchange (CME). Surprisingly, the empirical evidence shows that the inventory-based models, which suggest that the market makers' quotes and the inventory level are negatively correlated, cannot provide an accurate description of behaviors of market makers, since the evidence shows a strong and consistent positive correlations between reservation prices and inventory, i.e., market makers with long inventory positions are tend to increase their prices and sell aggressively.

Manaster and Mann (1996) identify two major weaknesses with their data-set which could change their results significantly. First, the CME's mutual offset system with the Singapore International Monetary Exchange (SIMEX) allows CME traders to unwind inventory positions on SIMEX overnight. Yet, the SIMEX trades are not included in their data-set. Second, it is impossible to recognize cases that locals trade on their own account through other

members. These are significant weaknesses given that locals normally trade through other members when they intend to modify the inventory position while attempting to conceal their activities from other investors. Thereofre, Frino et al. (1999) re-examine the predictions of inventory-based models with a new improved data-set provided by the Sydney Futures Exchange. In accordance with Manaster and Mann (1996), the authors also provide evidence that the predictions of inventory-based models are consistently contradicted by empirical results.

2.1.5.3   Information-based models

The inventory approach reviewed in the previous section provides some insights into the pricing behavior, among them, one important point is that the bid-ask spread is determined by transaction costs. Different from the inventory models, a new approach developed which models the market prices without relying on transaction costs, but on the influence of information. As argued by Glosten and Milgrom (1985), Kyle (1985), Easly and OHara (1987), even if no other costs incurred, the bid and ask spread can still arise due to information-based trading.

The origin of the information-based models is credited to a short but insightful paper by Bagehot (1971) with his distinction between market gains and trading gains. A *market gain* is arising from price changes and any dividend obtained from an asset during a given time period. Traders can either realize the market gain by holding the asset or trade with each other. The gain from the trading activity is called the *trading gain*. The concept of market gain is that most investors get profit when market price goes up, if the prices goes down, most investors face a loss. As prices move up and down over time, investors are expected to receive a market rate of return. The concept of trading gains, however, implies that the normal traders actually make a loss comparing with the market return as a result of the information costs.

According to Bagehot (1971), trading is a game of "zero-sum", in which one

trader to win, another must lose. Therefore, traders with special information tend to make a superior profit consistently at the expense of liquidity provider and other traders. Liquidity provider, or market maker, is assumed to face two types of traders in the market: informed traders and uninformed traders. *Informed trader* knows or better estimates the fundamental value of assets with some private information, *uninformed* or *liquidity traders*, on the other hand, make transactions with any private information.

Since informed traders know the "true" value of an asset, they will sell only if the market price is higher than its value and buy when it is lower. The market maker, as a result, will always suffer a loss from trading with these traders. In order to maintain the function of a market, the market maker's loss should be offset by making a gain from trading with uninformed traders. As the difficulty of distinguishing between informed and uninformed traders in an anonymous market, the market maker would charge a higher bid-ask spread to everyone for breaking even. The trading cost is the price that uninformed traders pay to compensate the losses of the market maker, which is also called *adverse selection costs*, i.e., the losses of uninformed traders to informed traders.

As the first to formalize the concept from Bagehot (1971), Copeland and Galai (1983) develops a simple one-period model of market maker's price setting in an asymmetric information system to study the information effects on the bid-ask spread. In this framework, a single risk neutral market maker trades with a population of investors. The stock price is a random variable, denoted P, taken from a known distribution F with the density function $f(P)$. Informed traders are assumed to know the actual value of the stock. Uninformed traders, who transact for exogenous reasons, know only the price process not the true value. In this model, the probability of any given trade with informed trader is π_I and $(1 - \pi_I)$ for uninformed trader.

Copeland and Galai (1983) also assume that the probability of an uninformed trader to buy is π_B, the probability to sell is π_S, and the probability of no trading is π_N. The informed trader make the transaction which maximize the

profit. The trade sizes are all fixed. And, the price-elastic demand functions of traders are also considered in the price setting process. When the dealer trades with informed, the expected loss for the market maker is

$$(2.7) \qquad \int_{P_A}^{\infty} (P - P_A) f(P) dP + \int_{0}^{P_B} (P_B - P) f(P) dP,$$

in which P_A and P_B are the dealer's ask and bid prices. On the other hand, if the trader is uninformed, the market maker's expected profit is measured as

$$(2.8) \qquad \pi_B (P_A - P) + \pi_S (P - P_B) + \pi_N (0).$$

The dealer does not know which type of investor he is trading with. Therefore, he weights his expected loss and gain by the probability of π_I and $(1-\pi_I)$ respectively. Therefore, $-\pi_I$ times 2.7 plus $(1 - \pi_I)$ times 2.8 provides the dealer's objective function. By maximizing the objective function, the optimal bid and ask prices (which are both positive) can then be derived. The results show that the size of bid-ask spread is influenced by various market parameters. It increases with higher price level and greater return variance; it is also inversely related to market activity, depth, and continuity, and finally it is negatively correlated with the degree of competition. Thus, a monopolistic dealer will be more likely to establish a wider spread than will perfectly competitive dealers.

2.1.6 Auction market

Different from dealer markets where traders buy assets at a dealer's ask price and sell at a dealer's bid price, in an auction market, investors trade directly with each other without relying on the dealer. In this market, buyers or sellers submit competitive bids or offers at the same time. The trading prices, which aggregated over a given period of time, represent the highest bid price that a buyer is willing to pay and the lowest ask price which a seller is willing

to sell at. Similar to the previous part, models applied in this market also examine the informational content of prices. Most models of the auction market in the market microstructure literature are Kyle type models which follow the study of Kyle (1985), Kyle (1986a), and Kyle (1986b). Due to the outstanding position, the work of Kyle (1985) is reviewed in this section.

A single asset, as an assumption of the particular model, is exchanged and traded among three groups of traders, including *insiders*, or risk neutral informed traders, who can access to a private information on the liquidation value of the asset before trading, *noise traders* who have no information and trade randomly based on exogenous needs for liquidity, and market makers who set prices with information they have on aggregate net order flow or order imbalance.

At each trading round, insiders and noise traders simultaneously submit market orders with the quantities they will trade to a market maker. While an insider's choice is made with information on liquidation value of the asset as well as the past prices and quantities he traded, noise trader places orders with random quantities independently from any information or previous trading record. After submitting orders, the market maker set a single trading price which is assumed to equal the expected asset value given the market maker's information consisting of the present and previous aggregate quantities traded by both insiders and noise traders. And, the aggregate quantities cannot be separated into individual quantities traded by insiders or noise traders due to the trading anonymous. Thus, the expected profit for a market maker is zero. The risk neutral insider, as an intertemporal monopolist, is assumed to maximize the objective function of expected profits by exploiting the monopoly power over time.

By applying a simple linear structure, Kyle (1985) investigates a unique "sequential auction equilibrium" that prices and quantities follow linear functions of trading observations. As the time interval between trading rounds reduces to zero, the author indicates that the discrete sequential equilibrium converges to a limit called continuous auction equilibrium. In both models,

the superior information of insiders is incorporated into market prices slowly at a constant rate. In addition, all of the private information in a continuous auction equilibrium will be reflected from the prices before the end of trading.

Kyle (1985) also defines the principal elements of a liquid market to be *tightness*: the narrowness of the bid-ask spreads; *depth*: the influence of execution price deterioration on the potential market order volume grows, and *resilient*: the speed of prices recover from a random shock.

Based on the of the Kyle models, several extensions have been made. For example, Bondarenko (2001) introduces a model with more than one market maker who quote prices to maximize their own profits with some strategies. As the number of market makers goes up to infinity, the results derived from the new model are found to be hold. With the continuous-time Kyle model, Back et al. (2000) denote that "the market would have been more informationally efficient has there been a monopolist informed trader instead of competing traders." However, they also notice that if comparatively large amount of private information remains near the end of trading, it would impair the market liquidity severely.

2.1.7   Limit order market

The need for trading liquidity is different across traders, but the two parties, buyers and sellers, have to trade simultaneously to clear orders in the market. This inconvenience is addressed in one of the three broadly observed markets, which are dealer markets, auction markets, and limit order markets. With the intervention of a market maker, dealer markets provide a desired liquidity to all traders at the same price. However, traders in auction markets need to either wait or trade prior to the desired time. In a limit order market, investors are allowed to decide immediacy with their choice of orders. Several papers, such as Rock (1988), Harris and Hasbrouck (1996), Seppi (1997),

Biais et al. (1999) and Foucault (1999), have studied the liquidity provision of limit order market by analyzing the limit order book.

The continuous limit order market is totally composed of limit and market orders. Market orders are executed at the best price available with a guarantee of immediate clearance. Limit orders, on the contrary, are only executed at a price meeting the requirements. A trader submitting a limit order sets a maximum (minimum) price at which he is willing to buy (sell). Though no immediacy and certainty of trading are provided, limit orders can help traders to have a better execution price compare to the market order price. Thus, the trade-off between the waiting costs and cost of immediacy, first suggested by Demsetz (1968), becomes dominant on deciding the order type.

In an early paper Cohen et al. (1981) derive a "gravitational pull" model to study a risk-neutral trader's preference in limit orders and market orders. The authors explain that more traders prefer to place market orders to ensure the order execution rather than limit orders when bid-ask spread narrows, as for the reason that the price benefits available to limit order traders are less attractive. The change of order submission strategy would, however, widen the spread, which in turn increases the potential benefits to limit order traders. Therefore, when choosing between limit orders and market orders, traders need to consider a dynamic balancing of the relative costs of execution uncertainty and costs of price improvement. Some other studies, including Kumar and Seppi (1994), Parlour (1998), Harris (1998), and Foucault (1999), indicate that limit orders are more attractive than market orders as the spread increases.

With dynamic models, Parlour (1998) studies how the order submission decision is affected by the depth available at the inside quotes. It is shown that depth on either side of the trading affects on order placement decisions. An arriving selling order is more likely to be a market order if the size is very large on the sell side, since a new limit sell order would be waiting at the end of a long queue with a low probability of trading. Moreover, if the quantity of purchasing is large, a new arriving ask order is more likely to be limit order.

In accordance with this theoretical study, Biais et al. (1995) demonstrate that traders are more willing to submit limit orders within the quotes when the depth at the quotes is large. A resent study by Ellul et al. (2002) analyze limit order submission strategies using data from New York Stock Exchange. Their result is mostly consistent with Parlour (1998) suggestion that depth on each side of the order book and the bid-ask spread are determinant.

The study of Foucault et al. (2005) investigate the influence of investors' impatience on bid-ask spread, order submission strategies and market resiliency. The result shows that the limit order book dynamics are driven by the proportion of patient investors and the order arrival rate. With a large proportion of patient traders or a low order arrival rate, traders turn to place more limit orders, and the market is therefore more resilient. In addition, a reduction in the tick size is found to increase the average bid-ask spread, but reduce the market resiliency.

To conclude, we provide a short review on market microstructure theory in this section. In the next section, some models using zero-intelligence traders applied to simulate the financial market are presented.

2.2   Market simulation

The traditional finance paradigm studies financial markets using models based on the rational expectations hypothesis: traders in the markets are assumed to be rational representative agents who continuously maximize their expected utility in an inter-temporal framework without leaving any profit unexploited. Traders compute the fundamental value of assets and price them accordingly, and immediately incorporate new information in the asset valuation. This conventional framework is appealing from a market-level perspective, however, it necessitate an unrealistic burden on investor behavior. In addition, it has been found that traditional framework based on the rational expectations hypothesis fail to understand basis facts about

the aggregate stock market and individual trading behavior (Barberis and Thaler (2002)), and are unable to reproduce the stylized facts (Gaunersdorfer and Hommes (2001) and Carceles-Poveda and Giannitsarou (2008)). Stylized facts are statistical properties of financial time series, which are robust across many instruments, markets and time periods. The common stylized facts include absence of autocorrelation of returns, volatility clustering, fat-tailed distribution of returns, and multiscaling, see Cont (2001), Anirban et al. (2011), and Sewell (2011). As mentioned by Sewell (2011), these general properties are so consistent that they are considered as truth.

Different from the traditional paradigm, in modern behavioral finance, some economists, such as Irving Fisher, Maynard Keynes and Harry Markowitz[9], argue that prices could be affected by individual psychology, and some financial phenomena could be better analyzed using models in which traders are not fully rational. As stated by Hirshleifer (2001), "The basic paradigm of asset pricing is in vibrant flux. The purely rational approach is being subsumed by a broader approach based on the psychology of investors". Other research, for example, Daniel et al. (1998), Cross et al. (2005), and Lamba and Seaman (2006) investigate the security prices with individual psychology. However, it is found that this approach is also not able to replicate the stylized facts adequately.

During the last few years, literature on financial markets has been focusing on the influence of institutional structures on market dynamics, with the consideration of traders' heterogeneity and complex interactions, see Calamia (1999). Heterogeneity could involve risk aversion, asset preferences, demand for liquidity, information and learning process; while the interaction between the heterogeneous behavior and institutional features could also affect important variables, for instance, the trading volume, transaction costs, price volatility, information processing and so forth. Therefore, they have become very crucial issues in the actual markets.

Decentralized financial markets consist of huge amounts of agents involved

---

[9] See, respectively, Fisher (1930), Keynes (1936), and Markowitz (1952)

in local repetitive competitive interactions which give rise to trading protocols. A large number of models have been developed to simulate the financial market and reproduce the observed stylized facts through the interactions of individual traders. Although many models based on the rational expectations hypothesis have good predictive power, it is argued that traders are unlikely to exhibit rational behavior in real world as they may lack necessary information or cognition. Over the last two decades the rational representative agent analytical model of the economy has been gradually replaced by a bounded rational heterogeneous agent computationally oriented evolutionary framework, so-called agent-based modeling, which is a powerful individual based computational simulation technique.

This section presents a brief overview of models used to study financial markets, including behavioral finance models, agent-based models, models with zero-intelligence agents and zero-intelligence-plus agents.

2.2.1   Behavioral Finance

Behavioral finance, which combines individual behavior and market phenomena, examine the impacts of social, cognitive and emotional factors on psychological decision process and the subsequent consequences for trading prices, returns and anomalies in financial markets. As defined by Frankfurter and McGoun (2010),

> Behavioral finance, as a part of behavioral economics, is that branch of finance that, with the help of theories from other behavioral sciences, particularly psychology and sociology, tries to discover and explain phenomena inconsistent with the paradigm of the expected utility of wealth and narrowly defined rational behavior. Behavioral economics is mostly experimental, using research methods that are rarely applied in the traditional, mainstream finance literature.

There are a large number of books and surveys on behavioral finance, see Shefrin (2000), Goldberg and von Nitzsch (2001), Barberis and Thaler (2003), Camerer et al. (2004), Ritter (2003), Altman (2006), Peterson (2007), Bruce (2010), Ackert and Deaves (2010), and Lovric (2011).

Behavioral finance has two building blocks: limits to arbitrage and cognitive psychology (Shleifer and Summers (1990), Barberis and Thaler (2003)). The limits to arbitrage[10] discusses that rational investors may not be able to exploit opportunities generated by less rational investors. In the traditional framework where traders are rational and no frictions are considered, a security's price equals its fundamental value, which is the discounted sum of expected future cash flows. To obtain the expectations, investors accurately process all available information. This is known as the the Efficient Market Hypothesis (EMH), which suggests that market prices fully reflect all available information. As soon as there is a deviation from fundamental value, or a mispricing, rational investors will immediately spot the arbitrage opportunity and make a positive profit with no risk, which in turn eliminate all the arbitrage possibilities by correcting the mispricing through the process of arbitrage. Thus, in an efficient market, no investment strategy can consistently achieve returns in excess of average market returns on a risk-adjusted basis. Behavioral finance, however, argues that when a mispricing, which are created by irrational investors, emerges, it may be difficult for rational traders to pursue the arbitrage strategies, due to the high cost, high risk and various constraints that involve in real financial markets. Therefore, rational traders may be unwilling to take arbitrage strategies, allowing the mispricing to remain for a long period of time (Barberis and Thaler (2003) and Lovric (2011)).

Cognitive psychology, the second building block of behavioral finance, refers to how investors think. Much of the behavioral finance literature focuses on individual investor psychology, and investigate how they aggregate across individuals, as well as its consequences for both investor performance and the

[10] For detailed discussion of the theoretical and empirical study of limits to arbitrage, see Shefrin (2000).

market dynamics. As early as 1960s, psychologists began to study economic decisions, e.g., Oregon Psychology Professor Paul Slovic studied the investment process from a behavioral perspective in 1969 (Slovic (1969)). And, a number of possible cognitive biases or deviations from rationality have been widely discussed in behavioral finance literature. For instance, heuristics are one of the most important concepts in behavioral finance. They are rules or strategies that can be used to solve various problems, but they can sometimes generate suboptimal solutions. The $1/N$ rule[11] is a simple heuristic that many people use it when deciding the allocation of their retirement money on various investment funds. Representativeness is another cognitive bias which means investors are likely to put more weight on recent experience rather than long-term averages. Moreover, it has been found that people are always overconfident when making decisions, they are also loss aversion and tend to avoid the pain and regret of having made a bad investment. For a detailed review of the heuristics and biases, see Barberis and Thaler (2003), Ritter (2003), and Lovric (2011).

Behavioral finance is a rapidly growing area that studies the influence of psychology on financial decision making and financial markets. Bondt and Thaler (1985) realize that investors are likely to overreact to unexpected and dramatic news events, which leads to substantial weak-form inefficiencies in stock markets. Daniel et al. (1998) examine the market under- and overreactions with the proposed theory of securities market based on investor overconfidence, and variations in confidence resulting from biased self-attribution. They discover that overconfidence causes negative long-lag autocorrelations, excess volatility. Biased self-attribution, on the other hand, results in positive short-lag autocorrelations, short-run earnings drift, and negative correlation between future returns and long-term past stock market and accounting performance. Barberis et al. (1998) also use concepts from psychology to study the empirical findings of overreaction and underreaction. They find that securities with a long record of good news are likely to become overprices and have lower average return subsequently. Hirshleifer (2001) provides a uni-

---

[11] To invest equally across all N funds under consideration.

fied explanation for most known judgement and decision biases on the basis of investor psychology and security pricing. Barber and Odean (2001) find that men trade more and get less profit on average because men are more overconfident than women. People with more confidence tend to trade more and, due to transaction costs, earn lower returns.

There are, however, several criticisms of behavioral finance. As stated by Barberis and Thaler (2002), "In some behavioral finance models, agents fail to update their beliefs correctly. In other models, agents apply Bayes' law properly but make choices that are normatively questionable, in that they are incompatible with SEU[12]." Another criticism is that by using which bias to emphasize, one can predict either under- or overreaction. As explained by Ritter (2003), "one can find a story to fit the facts to ex post explain some puzzling phenomenon. But how does one make ex ante predictions about which biases will dominate?" Some also argue that even if a small number of investors behave irrationally, the majority of rational investors will prevent security prices from shifting too much from the intrinsic value, Schindler (2007).

2.2.2 Agent-based models of financial markets

A novel bottom-up approach to investigate stock markets comes from the area of computational finance as artificial financial markets. Agent-based artificial markets, which always consists of a number of heterogeneous and boundedly rational traders, can be applied to study trading behavior, the influence of market mechanism, price discovery process, and the stylized facts of real-world financial time-series. The first modern agent-based model in finance is developed by Kim and Markowitz (1989) who study the stock market crash in 1987. As the emergence of new information could not explain the dramatic decrease, Kim and Markowitz (1989) attempt to find out other determinants of stock price volatility, such as hedging strategies and portfolio insurance. In

---

[12] SEU refers to Subjective Expected Utility.

this agent-based simulation, they explore the relationship between the agents with portfolio insurance strategies and market volatility. The results show that the market performance varies dramatically if the fraction of one type of trader changes relative to another. A market with more portfolio insurance related investors would be much more unstable. It is also found that trading volume, price volatility, and the size of extraordinary price changes is rising when the proportion of portfolio insurance investors increase.

More recently, Leigh Tesfatsion provides a large amount of literature on agent-based computational economics (ACE)[13], which is the portal to the application of agent-based model in computational finance and economics. In Tesfatsion (2002), she presents that ACE is a research framework that relying on the power of computers to explore the evolution of decentralized market economies with controlled experimental conditions. In the first place, ACE modelers construct a financial market with a population of agents, and specify the initial conditions of the market and the attributions of agents such as behavioral norms and internally stored information. Then, the market evolves over time without any further input or intervention from modelers. All events occur during the time period are caused by the historical time-line of agent interactions.

Another leading researcher in the field of agent-based model in finance, LeBaron (2001) reviews several design issues on building artificial agent-based financial markets. For the most important design question of agents design, LeBaron presents three types of modeling. According to him, " "The simplest and most direct route is to model agents as well-defined dynamic trading rules modeled more or less as strategies used in the real world." Using this method, insights about the interactions and tractable precise results are obtained with fewer costs. However, some criticize that markets may not follow these trading strategies continuously without any modification, and agents in these markets should trade with well-defined objective functions. The second way of modeling agents is to assume that they behave randomly

---

[13] See, for example, Tesfatsion (2001), Tesfatsion (2002), Tesfatsion (2003), and Tesfatsion (2006).

as zero-intelligence subject to a budget constraint. Finally, LeBaron mentions that learning and adapting strategies as in markets, for example, the Santa Fe artificial market[14], could be applied to model agents.

Schulenburg and Ross (2000) explore the reliability of artificially intelligent agents when applied to real markets with a simple adaptive agent based model. They discover that artificial agents, by both adaptive and non-adaptive simulations over ten year's real stock market data, are able to design successful market strategies to outperform the basic strategies such as buy-and-hold or saving money in a bank.

As indicated by Hommes (2005), interpretations of stylized facts observed in the financial market data provide an important motivation for agent-based models. Stylized facts are a set of properties, such as absence of autocorrelation, fat tails, volatility clustering, etc.[15], that have been widely observed across various instruments, markets and time periods. A large amount of literature investigates the stylized facts in financial markets with agent-based models, such as Hommes (2002), Cincotti and Pastore (2005), and Iori (2002). LeBaron et al. (1999) analyze the time series properties of asset returns using an experimental simulation with artificial intelligent traders who predict about the future risk and return to make decisions on buying or selling stocks. The simulated market and model are able to replicate several real-market features, which include volatility persistence, fundamental and technical predictability and leptokurtosis. A common way to quantify the divergence from the normal distribution is using the kurtosis of the distribution. The kurtosis equals three for a normal distribution, while a positive value of excess kurtosis indicates the existence of heavy tail. In this research, they find that the agent behavior is consistent with these properties.

Lux and Marchesi (2000) investigate three uni-variate properties of the financial assets, unit roots in levels, fat tails of the distribution and volatility

---

[14] The Santa Fe artificial market is a computer-based model with heterogeneously learning traders, see Arthur et al. (1997), Lebaron (2002), and Ehrentreich (2007).

[15] See Cont (2001) for a set of stylized facts that are commonly observed in various types of financial markets.

clustering. Using an agent-based model, they find these stylized facts in the artificial market, and present an explanation of the volatility clustering. Both chartist and fundamentalist strategies are considered in the research where traders are free from switching between the behavioural modes based on the observed payoffs. The result shows an efficient and stable market at most of times. Clustered volatility occurs if the proportion of chartist traders surpasses a threshold value. And, it is explained that occasional temporary instability of the market results in the volatility clustering.

Rabertoa et al. (2001) develop a simulated agent-based model in an attempt to perform computational experiments with different types of artificial traders. In the model, agents are endowed with finite resources in terms of cash and portfolio of assets. They randomly buy and sell in each period subject to limited resources, clustering and the volatility of previous periods. The developed model is capable to exhibit key properties of financial time series, i.e. volatility clustering and fat tails.

To construct a more sophisticated model that is consistent with many observed stylized facts in financial markets, Shimokawa et al. (2007) propose an agent-based equilibrium model with loss-averse features of traders, which is based on psychological study on decision making under uncertainty. Prices generated through the model are consistent with the equity premium puzzle, volatility clustering, asymmetry of distribution of returns, and cross-correlation between return volatility and trading volume. The model also shows that the market liquidity emphasizes these properties. Finally, they summarize the effects of factors forming equilibrium prices are loss-aversion parameter, the precision of private information and the reference point selection. An increase in loss-aversion and a drop of information precision magnify the price distortions, and the price properties, especially the strength of autocorrelations, are crucially influenced by the reference point selection.

In a recent paper, Martinez-Jaramillo and Tsang (2009) modeled a very simple financial market with sophisticated agents to investigate the co-evolution of the agents and the impacts of their strategic behavior on the price. The au-

thors also examine under what conditions the endogenously generated price would reproduce the statistical properties of real prices. Three different types of traders are studied, including technical, fundamental and noise traders. Technical traders, who believe that all the market information is already reflected in the price movement which shows patterns and trends, use previous trading information and mathematical indicators to make the trading decisions. Fundamental traders, on the other hand, attempt to evaluate a security's intrinsic value by studying related economic, financial and other factors. If the price departs from the intrinsic value, the fundamental traders will decide what position to take with that security, either buy or sell. Noise traders, whose decisions to buy, sell, or hold are irrational and erratic, are included in the research to represent a source of noise. The experimental results indicate that when learning happens, it helps not only in improving the traders' wealth, but also for producing more closely the stylized facts. In addition, heterogeneity has also been found important to reproduce the statistical properties of the price and the returns.

2.2.3   Zero-intelligence

It has become popular in recent years to use traders following very simple trading protocols that do not resemble rational behavior as assumed in most models in market microstructure theory as well as mechanism design. The zero-intelligence and zero-intelligence-plus traders are extreme forms of traders used in these models by assuming nearly random behaviour. The use of the simple algorithmic trader allows to investigate properties which are attributable to various forms of market mechanisms, regardless any impact from the traders' behavior. It also helps to replicate realistic statistical properties of prices which are otherwise very elusive in models of financial markets as well as the modeling and explanation of crashes in stock markets.

Smith (1962) reports on experiments performed in his simple classroom markets after the double auction of stock and commodity exchanges, which consist of a group of human traders who are entitled to buy or sell with a

limit price. His research is one of the first to indicate that small groups of traders can rapidly approach the rational equilibrium price. At the same time, Becker (1962) shows that agents' random choice behavior and budget constraint are sufficient to generate several basic economic features such as the proper slopes of demand and supply curves.

Traders in Smith's experiment are motivated by seeking 100 percent of the maximum exploitable profits; Becker assumes that the equilibrium results from supply and demand function are yielded through the traditional Walrasian tatonnement mechanism. To synthesize the generalizations by Smith and Becker, Gode and Sunder (1993) run a computer experiment to examine the market behavior with both constrained and unconstrained zero-intelligence traders (represented by ZI-C and ZI-U) who do not learn, but are almost random in the trading behavior. ZI-C traders are subject to the budget constraint, i.e., the money-losing transactions are not permitted. On the contrary, ZI-U traders are free from the budget constraint, and they are allowed to place a bid above their redemption value or an ask below the cost. The authors demonstrate that a non-Walrasian mechanism, a double auction market, can be allocationally efficient even traders without any sophisticated strategies to seek profit. In a double auction market, buyers can submit bids or raise existing bids; while sellers can enter offers or lower existing offer. Transactions take place at a match or cross of bid and offer.

Three notable features of the ZI-C trading behavior are found. First of all, there is no learning from day to day as expected from traders who have no memory. And, the volatility of ZI-C trading prices is lying between the stable prices set by human traders and the high volatility of ZI-U prices. Evidence shows that the presence of a budget constraint tends to move the trading behavior of ZI towards human behavior. Finally, the ZI-C prices are found to converge slowly to the equilibrium level during each trading day.

In Gode and Sunder (1993) the market discipline imposed on traders is largely responsible for achieving high levels of allocative efficiency, and it is also an important determination in the convergence of the transaction prices to

the theoretical equilibrium level, regardless of the intelligence, motivation or learning of the traders in the market, as the absence of rationality and motivation tend to be offset by the market structure.

Another research applying zero-intelligence agents is provided by Cincotti et al. (2005). This paper presents a multi-assets artificial financial market characterized by different kinds of stocks. With limited resources, zero-intelligence traders are designed to follow a random allocation strategy. Within the framework, several stylized facts are recovered in this paper, such as the volatility clustering, reversion to the mean, and fat-tail distribution of returns.

In another research, Krause (2006) explores return properties of a limit order market using ""near-zero-intelligence" agents, who do rely on calculations with a minimal set of trading rules. Each trader can only hold one unit of the asset and submit limit orders for one unit. No budget constraint is included, but the restriction for holding only one unit of asset could be viewed as a soft budget constraint limiting the demand. The limit prices of buy orders are random and following the distribution $lnB_i \sim iidN(\mu_t, \sigma_{buy}^2)$, where μ_t is the average asset price in long term. The limit prices of sell orders are also randomly taken from the distribution $lnS_i \sim iidN(lnP_i, \sigma_{sell}^2)$, which indicates the selling price depends on the price P_i the trader paid for in his last buy. Transactions take place at a price when the trading volume is maximal. After the order is matched and cleared in the market, a buyer switches to a seller and vice versa. If orders are not cleared after a long period, they are replaced by new limit orders with the same submission way.

The result of Krause (2006) shows that minimal intelligent traders, submitting random limit orders in the double auction market, are capable to generate the stylized facts of fat tail distribution of returns and multi-scaling. However, they failed to replicate the long term memory of volatility and the correlation between trading volume and volatility.

2.2.4 Zero-intelligence-plus

Cliff and Bruten (1997) review the early experimental work of Smith (1962) and Gode and Sunder (1993) and criticize their central argument that the convergence of transaction price to the equilibrium price is primarily determined by the market structure. If this is true, they state that developing markets for trading agents would be irrelevant. However, this research demonstrates that the average trading prices of zero-intelligence traders are only close to the equilibrium level in very specific circumstances when the supply and demand curves are symmetric, i.e. the magnitude of the gradient is roughly equal. Furthermore, they replicate Gode and Sunder (1993) observation of convergence of trading prices towards the equilibrium price within a trading day in two markets in which the supply and demand curves are symmetric or the supply curve is flat, but such convergence is not observed in the "box design" markets where the supply and demand curves are both flat with either excess supply or excess demand.

From Cliff and Bruten (1997) it is apparent that zero intelligence is not enough to account for the equilibrium convergence, some " "intelligence" in the form of sensitivity to current and previous market events is necessary. Consequently, Cliff and Bruten further investigate and introduce simple trading agents with adaptive capabilities, so-called zero-intelligence-plus (ZIP) traders, as a solution to the pathological failures of ZI traders. On the basis of ZI traders, ZIP traders, with the employment of elementary machine learning, systematically change their current bids and offers based on the information about the bid and offer levels accepted in the last round of trading. Experimental result shows that interacting of ZIP traders is similar to those used in Smith (1962) and Gode and Sunder (1993), but the performance of ZIP traders is significantly closer to the human experimental data than the performance of ZI-C traders in terms of equilibration and profit dispersion.

Following the introduction of ZIP strategy, Das et al. (2001) examine the performance of ZIP agents and present the result that the ZIP traders consistently outperforms human traders in a series of laboratory experiments.

The ZIP traders are found to "consistently obtain significantly larger gains from trade than their human counterparts."

2.3 Market mechanism design

2.3.1 Introduction

Mechanism design, as a sub-field of game theory and microeconomics, addresses the problem of decision making in distributed systems involving various self-interested agents with private information about preferences and capabilities. In a mechanism design problem, the essential goal of the designer is to find good system-wide solutions to distribute resources efficiently, i.e., to choose the optimal rules of the market so that the designer's objectives are achieved when market participants play their best strategies (see Parkes (2008) and Phelps et al. (2010)).

In recent years, market design has significantly broadened the scope of its application in computer science and operations studies; new attention has been given to the design of electricity markets (Bower et al. (2001), Sun and Tesfatsion (2007)), and online markets (Cliff (2003), Duffy and Unver (2008)); in solving the problem of distributed planning (Hunsberger and Grosz (2000), Bererton et al. (2003)), and the problem of combinatorial resource allocation (de Vries and Vohra (2003), Cramton et al. (2006)). Economists have investigate similar design problems in the context of auction theory (see Klemperer (2004) and Krishna (2002)) and mechanism design (see Varian (1995) and Jackson (2003).

Auction theory can be considered as an alternative approach to general equilibrium theory, in which an advanced and sophisticated micro-model of the marketplace can be built. As illustrated by Phelps et al. (2005), "Auction-theorists typically analyze a proposed market institution by defining a set of design objectives, and then proceed to show that these design objectives are

brought about when rational agents follow their best strategies according to a game-theoretic analysis. The task of choosing the rules of the market institution so that these objectives are brought about is called mechanism design.". However, there is no uniform standard by which we judge the "optimality" of an outcome. Instead, market designers will make the decision based on various settings. For example, in a market of assets auctioning, two different criteria can be applied for judgement: whether the revenue is maximized; or whether the allocation is efficient. The typical design objectives include, for example, allocative efficiency, budget balance, incentive compatibility, etc.

A large number of studies related to mechanism design have been emerging. For instance, Phelps et al. (2002), Phelps et al. (2003) and Cliff (2001) employ evolutionary algorithms to the mechanism design of double-auction markets. Byde (2002) investigates the auction mechanism design by using evolutionary game theory. David et al. (2005) study the optimal auction design for an English auction with discrete bid levels. Other studies, such as Mas-Colell et al. (1995), Marks (2006) and Phelps et al. (2010), provide accessible survey of mechanism design.

This section is to provide a brief overview of mechanism design. Auctions, as one of the simplest mechanisms to allocate resource and the most common types of financial markets, are the most popular market forms used when examining a mechanism design problem. Thus, in section 2.3.2 we provide a brief review of the auction theory and auction design. In the next section we discuss a number of simulation research on mechanism design, in particular with the application of zero-intelligence-plus strategies, agent-based models, and evolutionary search techniques.

2.3.2   Auction and auction design

As explained by Krishna (2002), "The word *auction* itself is derived from the Latin *augēre*, which means 'to increase' (or 'augment'), via the participle

auctus ('increasing')". Economists use the word "auction" to refer to a market institution in which resource allocation and trading price are determined by a set of explicit rules through the interaction between buyers and sellers. Auctions have been used to exchange a diversity of objects since ancient times as early as 500 BC, as recorded by Herodotus.[16] The Babylonians auctioned women for marriage annually; in the ancient Greek slave were auctioned; the Romans auctioned to liquidate the debtors' property which had been confiscated. In modern times auctions have been publicly and privately used to conduct a mass of economic transactions in numerous kinds of contexts, ranging from agricultural produce, livestock, tobacco, and fresh flowers to large tracts of vacant land, houses, art and antiques. In addition, governments become very keen on using auctions to sell everything such as treasury bills, oil drilling rights and other assets such as firms to be privatized. Today, the development of the internet, however, have greatly expanded the range of items sold by auctions, and raised the use of auctions with a great range of auctioneers. Almost anything can be traded by means of auctions. During the past decade, new auction markets have been set up to trade some special assets, such as transport, energy, and pollution permits, among which the mobile phone licenses auctions are the most famous new auction markets.

Auctions have become extremely important, not only because of its extensive use of allocating economic resources, but also the increasingly realized importance of auction theory, which has become the basis of much fundamental theoretical study, see Klemperer (2004). The auction theory investigates how traders act in an auction market and the properties of auction markets. It provides insights into many practical issues. The auction theory had been increasingly applied in practice since 1993, with the initial design and operation of the radio spectrum auctions in the United States, and the most notably one-hundred-billion dollar mobile-phone license auctions designed in 2000s. Another reason of the widely recognized importance of auction theory stated by Klemperer (2004) is that "by carefully analyzing very simple trading models, auction theory is developing the fundamental building

---

[16] See Krishna (2002)

blocks for our understanding of more complex environments". Moreover, Klemperer states that "auctions provide a very valuable testing-ground for economic theory-especially of game theory with incomplete information...and there has also been an upsurge of interest in experimental work on auctions." And, Milgrom (2004) demonstrates that "the qualitative predictions of auction theory have been strikingly successful in explaining patterns of bidding for oil and gas and have fared well in other empirical studies as well."

This section briefly considers some elementary auction theory, including the common types of auction, basic auctions models, and the celebrated Revenue Equivalence Theorem, followed by an overview of auction design, including electricity market design, mobile-phone license auction design, and on-line auction design.

2.3.2.1   Auction forms

Auctions can be classified in accordance with several distinct aspects. For instance, we can differentiate between open and sealed-bid auctions. All bids in the former auction form are open to the public, contrast to the latter one in which bids are not revealed publicly. In addition, we can also discriminate auctions based on the order submission process, i.e. either starting at a low price that increases progressively until no higher bids offer, or starting at a high price which then declines to a price that a bidder is willing to accept. Four standard auction forms that are widely used are described below, including the ascending-bid auction, the descending-bid auction, the first-price sealed-bid auction, and the second-price sealed-bid auction[17]. For simplicity, each auction type is interpreted based on a sale of a single product.

The *open ascending price auction*, also known as *English auction*, is the oldest and arguably the most prevalent type of auction in use. In this auction, the seller may announce prices, and bidders may call out prices themselves, or

---

[17] See Krishna (2002), Klemperer (2004), and Menezes and Monteiro (2008) for an overview.

bids with the highest current quote may be submitted electronically. Participants bid freely against one another with a successively raising price. The auction ends when only one bidder remains and no other participants are willing to submit a higher bid. The product is then sold to the remaining bidder at the final price at which the second-last bidder quit.

The *open descending-bid auction* is also named "Dutch auction" since it is widely used in the flower auction in the Netherlands. This auction form works in a contrary way to the English auction. The auctioneer starts with a high asking price at which no bidder is willing to buy. The price is continuously lowered until one bidder is willing to accept the offer. The winning bidder purchases the product at the last announced price.

Another common type of auction is the *first-price sealed-bid auction.* In this form of auction all bidders submit sealed bid simultaneously without seeing other participants' bids. In contrast to the English auction, a bidder in the first-price sealed-bid auction is only allowed to place a single bid. The object is sold to the bidder with the highest bid. The outcome in this auction form is identical to the one in the Dutch auction.

Finally, the *second-price sealed-bid auction*, also known as *Vickrey auction*, is similar to the first-price sealed-bid auction. However, instead of paying the highest bid, the winning bidder pays the second highest bid. This system is much alike the proxy bidding system applied by eBay, that the winner purchase the product at the second highest price plus a bidding increment.

Apart from these prevalent forms of auction, there are some other auction types. For example, a *reverse auction* is an auction type with an inverse role of buyers and sellers, contrast to the ordinary auction. In a reverse auction, several sellers offer their goods or services for bidding, and compete for the price at which a buyer is willing to accept. Another well-known market form is the continuous double auction (CDA), in which groups of buyers, who announce increasing bid prices, and sellers, who announce decreasing ask prices, submit orders simultaneously and asynchronously; a buyer is free to accept an offer, and a seller is free to accept a bid at any time. The CDA

is of particular interest in much academic studies, since it is the basis of most of the major financial markets. In addition, a *Walrasian auction* is an auction in which an auctioneer takes bids from both buyers and sellers of multiple objects, and determines the price and adjusts it over time based on the demand and supply. Furthermore, an auction in which a buyer can purchase more than one unit of the items is called a *multi unit auction*. If each of the items are sold at the same price, this auction is referred as *uniform price auction*; otherwise, it is referred to *discriminatory price auction*.

2.3.2.2   Basic auction models

One of the inherent features of auctions is the presence of asymmetric information. For example, sellers are uncertain about the values that buyers are willing to accept. If the seller knew the buyers' information, he could simply offer the item to the buyer with the highest price at or just below the maximum amount the buyer is willing to take. Models of auctions are generally classified into either of the following three categories based on the knowledge of information in the auctions, see Krishna (2002) and Menezes and Monteiro (2008).

In the *private-value* model, the value of the object is only known to each bidder himself. This information is privately held that bidders are unsure about the values attached by other bidders. And, the valuation of a specific bidder would not be affected by the knowledge of other competing bidders' value.

On the contrary, in the *common-value* model, the actual value of the object is the same for every participant, but it is unknown at the time of bidding. However, each bidder may have different information on the object's value. An example is the lease of a stretch of land with unknown amount of oil underground. A bidder would appraise the value depends on his estimates of the amount of oil underground. The privately known "signals" may be obtained from a self-conducted test, or an estimation from an expert. The

bidder would change his evaluation if he gets other bidders' signals, different from the private-value model where other participants' private information has no effect on a bidder's valuation. Wilson (1967), Ortega-Reichert (1968), and Wilson (1969) were the first to study common-value model.

Finally, Milgrom and Weber (1982) introduced the affiliated value model, which includes both private-value and common-value models as special cases. It assumes that each bidder has a private information, and a bidder's valuation is affected by others' private information. For example, valuation of an artwork may depend on the personal opinion of the value as well as others' private information since it would impact on the resale value and how easily it would be resold in the future.

2.3.2.3   Revenue Equivalence Theorem

A very common question in the economic analysis of auctions is whether two distinct auction markets could generate the same expected sales revenue. If one auction could achieve a higher selling price, then this auction type is much preferred by the seller. One of the most celebrated findings of auction theory, found by Vickrey in 1961, is the *Revenue Equivalence Theorem*, which states that the four standard auction types yield the same expected revenue to participants under certain conditions. The literature following Vickrey study the robustness of his result to the introduction of another bidding setups, e.g., Myerson (1981) and Riley and Samuelson (1981) verify that Vickrey's finding applied very broadly, and generalize his results. (Klemperer, 2004, p 17) gives a complete description regarding the Revenue Equivalence Theorem that:

> "Assume each of a given number of risk-neutral potential buyers of an object has a privately known signal independently drawn from a common, strictly increasing, atomless distribution. Then any auction mechanism in which (i) the object always goes to the buyer with the highest signal, and (ii) any bidder with the

> lowest-feasible signal expects zero surplus, yields the same expected revenue (and results in each bidder making the same expected payment as a function of her signal)."

2.3.2.4 Auction design

With the development of auction theory, various modern auctions, including spectrum auctions, mobile-phone license auctions, electric power auctions, timber auctions, emission markets, on-line markets, and various asset auctions, have been designed using different approaches. This section reviews the design of three popular markets, namely the electricity markets, mobile-phone license auctions, and online markets.

The design of electricity market has become an interesting topic since the liberalization of energy markets. The privatization of state electric power system took place in the 1980s in Chile and the United Kingdom[18]. The trend of electricity deregulation spreads subsequently to other countries such as New Zealand, Australia, and the United States. The deregulated electricity markets can be viewed as auctions [19], in which infinitely divisible quantities of homogeneous items are traded.

There is a large amount of academic literature on electricity market design. In 1992, a pioneering work by Talukdara and Ramesh (1992) suggests to use software to control electricity generation under the rapidly changed operating environment. Rasanen et al. (1994), for the first time, develop an object-oriented model of electricity demand-side load to study the electricity market design. Subsequently, Hamalainen et al. (1997) generate an "agent-based modeling framework" to examine both demand and supply side of the electricity markets. Lane et al. (2000) extend the previous work by introducing learning process using a genetic algorithm to the adaptive agents in a discriminatory-price k-double auction. For other research, see Maifeld

---

[18] See Klemperer and Meyer (1989) and Littlechild (1995) for an earlier study of UK electricity markets.

[19] See von der Fehr and Harbord (1993).

and Sheble (1996), Richter et al. (1999), MacGill and Kaye (1999), Talukdar (2002), and Cau and Anderson (2002).

In addition, many economists are interested in studying the design of electricity markets. Bunn and Day (1998) develop an agent-based model to study the generators's behavior and examine whether implicit collusion leads to high prices. Nicolaisen et al. (2000) conduct an experimental and simulation-based investigation of a double auction electricity market to evaluate market power and efficiency relating to mechanism microstructure. Bower and Bunn (2000) study the wholesale electricity market in England and Wales, and examine the impact of different auction formats on market prices. In particular, daily versus hourly bidding, and uniform versus discriminatory pricing auction forms are considered. Generators are represented by autonomous and adaptive computerized agents. A simple reinforcement learning algorithm is used to attain the goal of simultaneously maximizing profits and approaching a target utilization rate of the power plant portfolio. The bidding strategies of agents are modified at each round depending on the performance of in the last round. For other studies on electricity market design, see Bower et al. (2001), Petrov and Sheble (2001), Bunn and Oliveira (2003), Cau (2003), Bin et al. (2004), and Sun and Tesfatsion (2007).

The auctions of the third generation (3G) mobile spectrum license were some of the biggest and most complicated in history. The mobile license auction generated lots of debate. Remarkable attention has been intrigued by the significant amounts of money involved, and the huge differences in the amounts paid across countries are of particular interest, see van Damme (2002), Binmore and Klemperer (2002), Borgers and Dustmann (2002), Klemperer (2004), Seifert and Ehrhart (2005), and Hoppe et al. (2006).

The 3G auctions in most European countries, except for Denmark, were of the "simultaneous multiple-round ascending" type, which had been introduced and widely employed in the US, see Milgrom (2004) for details. The United Kingdom ran the world's first 3G auction, which was later described as the "biggest auction ever" (Binmore and Klemperer (2002)). In the UK auction,

bidders were fully informed and the auction market was very transparent. Bidders, who place the highest bid, became the "current price bidder" in the next round; and bidders who were not the current highest bidder could either place a valid bid, ask for a total of three waivers, or withdraw from the auction. If there were multiple bidders having the same highest price, the "current price bidder" was then randomly selected among these bidders. The auction ended when there were no other new bids.

The German and Austrian auctions, which had the same complex design, were less transparent with limited information. Two consecutive auctions were involved in the design. As described by Jehiel et al. (2003): " The first auction allocated licenses together with so called 'duplex' or 'paired' spectrum frequencies. The second auction allocated paired spectrum that has not been sold at the first auction, together with additional 'unpaired' spectrum. Both auctions were of the 'simultaneous multiple-round ascending' type." In Italian auction, a maximum of five spectrum licenses were auctioned. In each round bidders were ranked based on their bids. Bidders who did not submit the five current highest bids would either leave the auction or improve their bids. The Denmark auction was the last one amongst the 2000-2001 European 3G auctions. To stimulate entry the Danish regulator adopted a sealed-bid auction, and those with the highest four bids obtained a license. The Danish auction was discriminatory as bidders paying their own bid instead of paying the highest losing bid. A number of studies have described the 3G Europe auctions in details. Jehiel and Moldovanu (2000), van Damme (2002), Jehiel et al. (2003), and Klemperer (2004) provide insightful review of the 2000-2001 European 3G license auctions. Seifert and Ehrhart (2005) conduct a laboratory experiment to compare the performance under the UK and Germany auctions in terms of revenues and bidder surplus.

Since the advent of the Internet, an increasing number of transactions are carried out using online auctions. Some popular internet auction sites, like eBay and Amazon, are emerging and having a great success. Various forms of auction mechanism have been implemented by the different web-based auction sites, among which, the most common and widely used formats stated

in Duffy and Unver (2008) are ascending-bid format, second-price format, and a hybrid of the ascending-bid English auction and the second-price sealed bid auction.

Different from the traditional auctions, online auctions allow bidders to bid anywhere in the world at any time with Internet. Some online auctions run many days, providing a good chance for more bidders to notice and enter the auction. Lucking-Reiley et al. (1999) and Hasker et al. (2004) express that the longer online auction runs, the more bidders are attracted with higher prices. As found by Lucking-Reiley et al. (1999), prices are 24% higher, on average, in 7-day online auctions than in short period auctions. However, auctions with longer duration raise the question of the optimal timing of the bid. An interesting phenomenon is that of "late" bidding, which is placed very close to the end of the auction, gives other bidders little time to respond.

Another design question is whether bids in online auctions should be sealed or open, i.e., whether a bidder has a chance to observe how others' bidding activities evolve, and reacts to their bidding strategies during the auction process. Ivanova-Stenzel and Salmon (2004) show that in a private-value environment open bidding is much preferred by laboratory subjects. In addition, Ariely et al. (2005) conduct a laboratory human subject experiment which shows that the speed of learning the optimal bidding strategy is much faster in open auctions than in sealed-bid auctions. However, there are some other argument against the superior of open format. Cramton (1998) reports that open bidding fails to generate high revenues when bidders are risk-averse and competition is low. Moreover, Klemperer (2004) finds that collusion among bidders in open auctions may reduce the revenue.

2.3.2.5 Limitations of classic mechanism design

The traditional mechanism design and auction theory produce clear-cut results in many situations. However, the underlying assumptions of the theory may not be fulfilled in many real-world situations. The traditional approach

of mechanism design seeks closed-form solutions to evaluate the performance of economic institutions. As discussed by Marks (2006), under this approach, "Economic actors have been assumed to be perfectly rational, with the means to solve for equilibria outcomes in complex situations. Economists have sought to characterize the equilibria of economic interactions in terms of their existence, uniqueness, and stability, under this assumption. When the interactions among economic actors are strategic, the equilibria become Nash equilibria." In the real-world markets, however, not only the equilibrium characterization should be considered, but also the continual shocks which might hamper the system to approach the equilibrium. Marks (2006) lists four reasons of considering other techniques rather than the traditional approach. Firstly, it would be difficult to get solutions for the design of several markets, for instance, the continuous double auctions. And, the dynamic behavior or out-of-equilibrium behavior of an changing market environment cannot be characterized. Another reason is that "the assumption of perfect rationality and unlimited computational ability on the part of human traders is unrealistic". Economists tend to use game theory to describe the sophisticated behavior of the participants, which has been criticized for its 'hyper-rational' view of human behavior (see Varian (1995)). Finally, Marks (2006) suggests to include the learning process, since "a model without learning is not as realistic as one incorporating learning. Bunn and Oliveira (2003) note that many researchers [including Erev and Roth (1998)] have shown that learning models predict peoples behavior better than do Nash equilibria."

Apart from the above limitations, the ability of traditional market microstructure theory to provide insights into the optimal market structure or the mechanism design is limited. As each market structure requires a slightly different model to allow it being studied analytically, so far, no general framework exists, and comparisons between two or more market structures are limited. As noted by Kittsteiner and Ockenfels (2006), "One of the important lessons that can be learned from theoretical and empirical research is that there is generally no one-fits-it-all design. Every market environment is different and has its own particularities." Given the large number of possible market

structure, any optimization procedure involving different models would be unattainable.

Furthermore, when using different models for different market structures it is even more difficult to distinguish between the effect attributable to differences in the market structure and to different behaviors of traders. The results of conventional models are often dominated by the behavior of the agents and it proves difficult to distinguish their impact from that of the market structure. With classical market microstructure theory developed to analyze trading behavior in a given market structure, it is in its current form not suitable to be applied in optimizing market structures.

Moreover, Phelps et al. (2010) disclose another problem of using classical mechanism design to study the existing market institutions, such as the London Stock Exchange, which has been established far before the proposal of auction theory and game theory. Since trading rules applied in the market institutions may not be based on traditional auction theory, thereby, it would be difficult to analyze the mechanism design using the traditional approach.

In addition, most of the models in the current microstructure theory investigate the impact of a single or a small number of related trading rules on specific aspects of market performance, ignoring hybrid systems such as the specialist system of the NYSE which are much more difficult to model. And, in order to evaluate whether one market structure is superior to another it is necessary to use a measure of market performance. Unfortunately these does not exist a single measure which is generally accepted; depending on the context of the research the use of measure will vary. Apart from the problem of measuring market performance, the research on the optimal market structure is also hampered by the complexity of the models used, allowing only focusing on optimizing a single trading rule or a small set of closely related trading rules. Given that the trading rules are usually showing non-trivial interdependencies in their influence on the market performance, it cannot be derived from such isolated research how the optimal market structure will be found. Phelps et al. (2010) denotes that "In designing market

places...we often need to make such trade-offs between different objectives depending on the exact requirements and scenario at hand. We can often satisfactorily solve such multi-objective optimization problems, provided that we have some kind of quantitative assessment of each objective, yet classical auction-theory provides only a binary yes or no indication of whether each of its limited design objectives is achievable, making it extremely difficult to compare the different trade-offs."

To address these problems, a large number of design features or trading rules should be considered to provide a full picture of the rules auction markets should be designed to in a sufficient number of realistic circumstances. To represent the different interests of various market participants, the best way is to conduct the optimization of the trading rules evolutionary using a wide range of performance functions by employing multi-objective optimization. This allows us to determine the most appropriate sets of rules a market should consider without facing the need to balance the different interests of market participants. Finally, to investigate the market structure and its impact on the market performance, we need to eliminate the influence of trading behavior on the market performance. Thus, the traders in the market should be automaton traders following very simple trading protocols that do not resemble rational behavior. The use of appropriate automatons allows us to focus on the influence the market structure, i.e. set of trading rules, has on the outcomes.

2.3.3 Simulation techniques and experiments

As mentioned by Kittsteiner and Ockenfels (2006), "theoretical considerations alone typically do not guarantee a sensible market design. This is true in particular in those cases in which reality is too complex to be solved analytically...typically, theory needs to be complemented with other tools such as computational, empirical and experimental methods (see Roth (2002), Ariely et al. (2005))." Therefore, when conducting any adjustment on the market structure, it should be guided by theoretical as well as empirical research.

Research using simulation has been growing as computing techniques become more powerful. "In particular, laboratory experiments can be used to test the validity of various economic theories, to test hypotheses that cannot be investigated with the help of field data, and to test new market mechanisms that do not yet exist.", indicated by Kittsteiner and Ockenfels (2006). This section reviews the empirical simulation and experiments on the mechanism design with the application of Cliff's zero intelligence plus traders, agent-based models, and evolutionary search techniques.

2.3.3.1 Zero-Intelligence-Plus strategy in mechanism design

As discussed by Marks (2006), "Market performance may depend on the degree of 'intelligence' or 'rationality' of the agents buying and selling, which has led to computer experiments in which trading occurs between artificial agents of limited or bounded rationality". An alternative approach to evaluate market structure has been introduced by Gode and Sunder (1993). They developed a model in which traders submit random orders, the only constraint on their behavior being that they do not make a loss. Simulating the outcome of a market with such traders, they found that it was close to the equilibrium, despite traders not learning in their model. This very simplistic behavior of traders led them to be called zero-intelligent (ZI) traders and the resulting efficiency of the market is attributed to the market structure. As reviewed in section 2.2, Gode and Sunder (1993) indicate that "zero-intelligence agents can exhibit human-like behavior in CDA markets." However, Cliff and Bruten (1997) argue that "Gode and Sunder's result only holds in very specific circumstances and that, in general, some intelligence in the form of adaptivity or sensitivity to previous and current events in the market is necessary." Consequently, Cliff and Bruten (1997) propose a more sophisticated and robust algorithm refers to Zero-Intelligence-Plus traders (abbreviation ZIP) to solve the pathological failures of ZI traders. ZIP traders are simple autonomous artificial trading agents with adaptive capabilities to operate as buyers or sellers in auction market environments. These agents results in a

good performance on several measures of trading activity, such as allocative efficiency and profit dispersion.

Following the introduction of ZIP strategy, Das et al. (2001) examine the performance of ZIP agents and present the result that the ZIP traders consistently outperform human traders in a series of laboratory experiments. The ZIP traders are found to "consistently obtain significantly larger gains from trade than their human counterparts." Because of the favorable results and its usability in every market structure, ZIP traders have a good potential as trading agents in a mechanism design problem.

Although the use of computational simulated approaches for the agent decision problem has a relatively long history, the field of computational-mechanism design is an emerging area of research. As before it is very hard to design a new market, especially with a vast searching space of possible mechanism. Therefore, the application of *automated search* algorithms of the space of possible market types is very attractive. To explain how to design a market using computerized artificial agents, Walia (2002) states that

> "To automate the process of market design we need to parameterise the description of a given market type. Once we have a complete parametric description of a market mechanism, the design process can be looked at as a parameter optimisation problem with the optimum parameter values specifying the best design."

After the development of ZIP strategy, Dave Cliff began to make a line of research on the evolution design of auction markets mechanism using computer simulations, see Cliff (2001), Cliff (2002a), Cliff (2002b), Cliff (2003). In this section, we will briefly review the series of work by Dave Cliff, and some other related research with the application of ZIP algorithm.

Prior to the work of Cliff (2001), the auction mechanism in which traders participate has been specified in advance of any investigation in all market studies using artificial agents. Cliff (2001), for the first time, develops a

complete system for automated mechanism design. He uses the Genetic Algorithm (GA)[20] to optimize both the parameters of the ZIP agents, and the mechanism parameter, which is Q_s in the paper, denoting the probability that a seller will be chosen to make an offer in any given round of a continuous double auction. Once the Q_s is specified, the probability that a buyer will be chosen to make a bid, Q_b, can be easily obtained as $Q_b = 1 - Q_s$.

In a standard CDA with the assumption of only one quote occurs at a unique time, $Q_s = 0.5$, which means, at any time, the probability of the next order submitted by a seller is 50%. Similarly, the probability of receiving a quote from a buyer for the next shout is 0.5 as well. In this case, buyers and sellers have an equal opportunity to quote in CDA. In English auctions, $Q_s = 0$ since only buyers can quote while sellers stay silent in this framework; in Dutch auctions, on the other hand, $Q_s = 1$ as only sellers are able to quote. For other values, e.g., in a mechanism with $Q_s = 0.2$, it can be thought of that there will be 2 quotes from sellers and 8 quotes from buyers for every 10 shouts. Nevertheless, this type of market is never found in human-designed market, and it is not clear whether the *hybrid mechanisms* with value of Q_s other than $0, 0.5$ or 1 are superior to any other well-known market types.

Cliff (2001) proposes a model to define the market mechanism using the inventive parameter Q_s. The Q_s with values of 0, 0.5 and 1 are not regarded as three different forms of market in this paper, but rather considered as two end points and one midpoint in a continuous mechanism space. To explore the evolution of market mechanism, he defines the searching space, including the continuous double auction and one-sided auctions similar to the English auction and the Dutch auction. In spite of this, Cliff points out that "this space is continuously variable, allowing for any of an infinite number of peculiar hybrids of these auction types to be evolved, which have to known correlate in naturally occurring market mechanisms."

Cliff (2001) conducts three tests with different markets M1, M2 and M3. These markets have the same number of traders, but with different supply

---

[20] A genetic algorithm is an automated search algorithm, reviewed in section 3.2.1

and demand schedule. Through a GA search, the optimum value of Q_s for the evolved mechanisms are approximately 0.0001, 0.0714, and 0.1666, respectively. Cliff further compares the performance of evolved mechanisms with a standard CDA market. The results show that the evolved "hybrid" mechanisms outperform the normal CDA market. As described by Cliff:

> "While there is nothing to prevent the GA from settling on solutions that correspond to the known CDA auction type or the EA-like or DFA-like one-sided mechanisms, we repeatedly find that hybrid solutions are found to lead to the most desirable market dynamics. Although the hybrid mechanisms could easily be implemented in online electronic marketplaces, they have not been designed by humans; rather they are the product of evolutionary search through a continuous space of possible auction types. Thus, the results in this paper are the first demonstration that radically new market mechanisms for artificial traders may be designed by automatic means."

To extend the work of Cliff (2001), in which a single fixed market schedule is used to evaluate the optimum solutions, in a following paper, Cliff (2002a) modifies the supply and demand schedule by inflicting "shock changes" through the evaluation process, i.e., switching from one schedule to another. The *market shock* is imposed to reflect the situation where unexpected events occur in financial markets, leading to a dramatic change in trading behavior. The experimental result, again, reveals that new "hybrid" market mechanisms with $Q_s = 0.25$ is more efficient than the well-known CDA markets. Another evolutionary mechanism with the optimum value of $Q_s = 0.45$, which is close to the parameter value of a standard CDA, is found to be "no better but also no worse than the CDA". [21]

Subsequently, Robinson (2002) and Qin (2002) replicate and test Cliff's experiments. Both of them have found a similar results, and therefore confirmed

---

[21] A series of related studies of exploring continuous market space by genetic algorithms with ZIP traders has been conducted by Cliff, such as Cliff (2002a) and Cliff (2003).

the reliability of the results on evolving market mechanism. Preist and van Tol (2003) introduce a variant of Cliff's ZIP strategy which can be applied in persistent-shout auctions. Walia (2002) question the robustness of Cliff's results, with concerns of the exclusive use of the ZIP strategy in the work. They apply an alternative algorithm, a self-adaptive Evolutionary Strategy (ES), to investigate the space of possible mechanism types in a CDA populated by the zero-intelligence-constrained(ZIC) agents developed by Gode and Sunder (1993). With similar results, they confirm that hybrid markets are more desirable than the standard market forms, regardless of the trading algorithms used.

2.3.3.2   Agent-based models in mechanism design

As reviewed in section 2.2.4, zero-intelligence-plus (ZIP), proposed by Cliff and Bruten (1997), is one typical example of the agent-based models. Instead of using ZIP, the automated search algorithm has also been applied to various market mechanisms design problems with alternative agent-based models. As discussed by LeBaron (2005), agent-based models are well suited to mechanism design and market microstructure questions not merely because of a large amount of data, but also the desirable testing environment of market design with heterogeneous, adaptive strategies. "With several design trade-offs and the possible emergence of unforeseen performance in the system, agent-based analysis and design, in which the market system can be modeled as 'evolving systems of autonomous, interacting agents with learning capabilities' [Koesrindartoto and Tesfatsion (2004)], is increasingly employed.", stated by Marks (2006).

As discussed by Davidson and McArthur (2005), Sun and Tesfatsion (2007), and Weidlich and Veit (2008), a growing number of researchers have developed variety of agent-based models for the study of electricity market, see, for example, Bower and Bunn (2000), Nicolaisen et al. (2001), and Veit et al. (2006),

Nicolaisen et al. (2001) examines the exercise of market power and efficiency of a wholesale electricity market by using an agent-based computational model. They present a new auction mechanism, which is simulated applying experimental traders and reinforcement learning, a modified version of the Roth-Erev learning algorithm, referred to as MRE algorithm. This paper provides some insight into the sensitivity of market efficiency to the MRE reinforcement learning algorithm, which characterizes the learning behavior of the trading agents. They disclose that when trading agents refrain from bad judgement, market efficiency would be higher. However, the result may not be robust with the exercise of inappropriate learning process or bad judgement. Some other work, for instance, Koesrindartoto and Tesfatsion (2004), Koesrindartoto et al. (2005), and Sun and Tesfatsion (2007), investigate the potential use of agent-based computational methods for the study of the Wholesale Power Market Platform (WPMP), a complicated market design proposed by the U.S. Federal Energy Regulatory Commission as a widely adoption by U.S. wholesale power markets.

Hailu and Schilizzi (2004) conduct an agent-based computational experiment to study the performance of conservation auctions with the assumption that bidders are allowed to learn from previous experience. Bidders in this model are farmers or landholders bidding for environmental conservation contracts. On the other hand, a government agent with a fixed budget selects and awards contracts to the winning bidder. The results show that the efficiency advantages of the auction dissipate with repetition and learning, making the auction mechanism less attractive compare to simpler mechanisms like fixed payment schemes.

Chen and Chie (2008) set up an agent-based model to examine the evolutionary mechanism design of lottery markets with genetic algorithms. Two conclusions are made based on the results. First of all, the Laffer curve in economics is observed, which confirms the existence of an optimal lottery tax rate, in this case it is 40%. Secondly, the optimal tax rate from the simulation is very close to the empirical tax rate of 42 %, averaged over 25 lottery markets.

Hailu and Thoyer (2006) investigate the mechanisms design of multi-unit auctions using an agent-based model. They compare performance of three auction forms, including Vickrey, uniform, and discriminatory forms, based on two criteria, namely the budgetary outlay and social opportunity costs. The experimental results show that "the relative social cost efficiency of the three auction formats ... depends on the level of competition. For high levels of competition, the three auction formats perform identically. As competition declines, the efficiency of the discriminatory auction becomes inferior to the efficiency of the uniform and the Vickrey, which remain similar ... In terms of budgetary performance, the uniform and Vickrey auctions perform equivalently for high levels of competition ... the uniform auction induces slightly lower outlays than the Vickrey at high demand levels. More important, both loose their advantage over the discriminatory which becomes the least expensive procurement auction at high demand levels."

Marks (2006) first provides a lengthy survey of the emerging practice of mechanism design using both evolutionary and agent-based techniques. A large number of papers are reviewed, with special reference to the design of electricity markets and online markets using agent-based evolutionary techniques. Furthermore, he summarizes several possible design criteria or the objective to be optimized in the evolutionary process used in previous research. The possible criteria listed in the paper consist of maximizing seller revenue (one of the major criteria in the spectrum auctions design, see Milgrom (2004)); maximizing market allocative efficiency (which means the performance of the mechanism should be optimal under certain sense); discouraging collusion; discouraging predatory behavior; discouraging entry-deterring behavior; budget balance (the total payment made by traders equals zero, i.e., no money injected into or removed out from a system); individual rationality (the expected net gain to each trader in the mechanism should match at least the net gain of any other systems), and strategy-proofness ("participants should not be able to gain from non-truth-telling behavior", explained by Marks (2006)). Marks suggests that there are trade-offs between several design criteria applied in a computer simulation, which has been proved by

Myerson and Satterthwaite (1983), as "no double-sided auction mechanism with discriminatory pricing can be simultaneously efficient, budget-balanced, and individually rational." Marks (2006) explains it with a simple example: if a seller has an influential power, his revenue can be maximized, but at the loss of market efficiency, since the sum of buyer surplus and seller surplus is reduced. Therefore, the author suggests either to weigh all the criteria to derive a single objective, or to rank the market designs based on the remaining criterion that are all above a target level.

For other research in the area of market design applying agent-based models, see Bunn and Day (1998), Campbell et al. (1999), Hailu and Schilizzi (2004), Koersrindartoto (2005), and Choi et al. (2010).

2.3.3.3 Evolutionary mechanism design

Evolutionary computation, which uses ideas from natural evolution and learning, has been widely applied to the study of mechanism design. Evolution can be considered as a dynamic search through a continuous space of possible market designs to obtain the market structure that results in better performance. For example, the optimal tick size for an exchange to maximize the trading volume can be investigated using evolutionary computation. Evolutionary algorithms have two distinct features relative to other search algorithms. As explained by Sarker and Ray (2010), "they are all population based (consider a set of possible solution points in contrast to single solution in conventional optimization). Second, there is communications and information exchange among individuals in a population."

We have discussed in section 2.3.2.5 that the traditional approach of mechanism design fails to characterize the intermediate behavior of a dynamic interaction. In addition, when the search-space is very large, exhaustive search is very difficult to implement. Therefore, other techniques of learning and evolution in strategic environments should be applied for analysis of

mechanism design, which makes evolutionary algorithms superior than conventional techniques for studying the interaction between different strategies by systematically sampling the search space, see Phelps et al. (2010).

With the development of agent-based modeling and optimization techniques such as Genetic Algorithms, evolutionary simulation techniques have been widely applied in the mechanism design problems.[22] Cliff (2001) is the first to employ evolutionary computing to the double auction mechanism design problems, reviewed in section 2.3.3.1. Because of the successful operation of evolutionary algorithms employed in Cliff's experiments, use of evolutionary algorithms for mechanism design becomes prevalent in the related research. A similar study of exploring the evolution of economic market mechanisms is conducted by Phelps et al. (2002). With application of genetic programming (GP), they co-evolve agents' trading strategies to analyze the development of auction pricing rules. This is followed by Byde (2002), in which an evolutionary-based method is applied to examine the mechanisms of standard first-price and second-price sealed bid auctions. The result presented by Byde shows that a new form of sealed-bid auction mechanism is more superior than the standards auction forms, and is optimal for a range of realistic scenarios. Other papers, such as Cliff (2002a), Cliff (2003), Robinson (2002) and Qin (2002) Walia (2002), Chen and Chie (2008), apply the evolutionary search techniques for the mechanism design problems as well. These papers have been reviewed in the previous sections. In addition, Marks (2006) and Phelps et al. (2010) provide a good review of the evolutionary simulation of mechanism design.

---

[22] Chapter 4 presents the evolutionary optimization algorithms.

3. An overview of evolutionary optimization algorithms

The term evolutionary algorithms (EAs) refers to a class of stochastic optimization search methods that are inspired by biological evolution such as mutation, recombination, reproduction, and natural selection applied to optimize a specific system within certain constraints. Different from traditional optimization techniques, EAs search a "population" of solutions instead of a single point. Eiben et al. (2003) describe the evolutionary algorithm as follows:

> "Given a quality function to be maximised we can randomly create a set of candidate solutions, i.e., elements of the function's domain, and apply the quality function as an abstract fitness measure - the higher the better. Based on this fitness, some of the better candidates are chosen to seed the next generation by applying recombination and/or mutation to them. ... Executing recombination and mutation leads to a set of new candidates (the offspring) that compete - based on their fitness (and possibly age) - with the old ones for a place in the next generation. This process can be iterated until a candidate with sufficient quality (a solution) is found or a previously set computational limit is reached."

Many different evolutionary techniques have been proposed, most notably Genetic Algorithms (GA), Genetic Programming (GP), Evolutionary Pro-

gramming (EP), evolution strategies (ES), and Estimation of Distribution Algorithms (EDAs), with similar underlying concept behind all these algorithms: see Fogel (1994), Back et al. (1997) and Eiben et al. (2003) for an introductory survey of evolutionary algorithms, and the working principles of the various algorithms.

Evolutionary algorithms have been widely applied in financial research, see e.g. Chen (2002), Day (2006) and Fernandez-Blanco et al. (2008). A wide range of evolutionary algorithms have been used in this research area, including Genetic Algorithms (Yuret and Maza (1994), Cliff (2001) and Cliff (2003)); Genetic Programming (Li and Tsang (1999) and Potvin et al. (2004)); Grammatical Evolution (O'Neill et al. (2001), Brabazon and O'Neill (2003b) and Brabazon and O'Neill (2003a)); Simulated Annealing (Ingber (1996), Ingber (1999)), and the Population-based Incremental Learning (Alexandrova-Kabadjova et al. (2006) and Alexandrova-Kabadjova et al. (2011)), which is a type of Estimation of Distribution Algorithm recently developed by Baluja (1994). In this chapter, we will first present a brief review of single objective optimization techniques, in the context of estimation of distribution algorithms, genetic algorithms, competitive learning, population-based incremental learning, further development and application of PBIL, and the empirical comparisons of different optimization algorithms. Finally, algorithms for multi-objective optimization problems are reviewed in section 3.3.

3.1 Estimation of distribution algorithms

Estimation of distribution algorithms (EDAs), also known as the probabilistic model-building genetic algorithms (PMBGA), are a recently developed set of evolutionary optimization algorithms (see Pelikan et al. (1999) and Larranaga and Lozano (2001) for a brief review). EDAs, which are first introduced by Muhlenbein et al. (1996), maintain a set of candidate solutions throughout the search. These solutions are randomly generated at the first iteration, which are then evaluated based on the objective function. Solutions with

better function values are selected. After the selection, a probabilistic model of the chosen solutions is generated, and a set of new candidate solutions, which replace part of the old ones in the original population, is sampled based on the probability model. This process is continued until the optimum is achieved or the termination criterion is fulfilled.

There is no traditional crossover and mutation process in EDAs. EDAs replace them with sampling new elements and learning from the probability distribution of the best gene of the population at each generation. Through this repetitive process of constructing a new generation based on the distribution model, and a following modification of the distribution model, the quality of the population is improving over time. Since new candidate population is sampled with a probability distribution, Zhang (2003) argues that "EDAs are promising methods for capturing the structure of variable interactions, identifying and manipulating crucial building blocks, and hence efficiently solving hard optimization and search problems with interactions among the variables." In addition, EDAs can also be applied efficiently in some complex problems where important interactions exist between the variables of solutions. These problems are, however, difficult to be solved by some evolutionary algorithms that do not employ models, such as genetic algorithms.

Many EDA-like algorithms have been developed, which can be categorized into three groups according to the interdependencies between the variables of solutions: univariate, bivariate or multivariate models. The univariate models are the simplest way to estimate the distribution of selected solutions by assuming the variables are independent, examples of univariate algorithms include population-based incremental learning (PBIL) (Baluja (1994), and Baluja and Caruana (1995)), univariate marginal distribution algorithm (UMDA) (Muhlenbein et al. (1996)), compact genetic algorithm (cGA) (Harik et al. (1999)) and learning automata based estimation of distribution algorithm (LAEDA) (Rastegar and Meybodi (2005)). The bivariate algorithms, which do not have the assumption of independency between variables, consider some pairwise interactions, examples of such algorithm are

mutual information maximization for input clustering (MIMIC) (De Bonet et al. (1997)), bivariate marginal distribution algorithm (BMDA) (Pelikan and Muhlenbein (1999)) and combining optimizer with mutual information trees (COMIT) (Baluja and Davies (1997)). Algorithms that cover multivariate interactions include factorized distribution algorithm (FDA) (Mhlenbein and Mahnig (1999)), Bayesian optimization algorithm (BOA) (Pelikan et al. (2000) and Larranaga et al. (1999)) and extended compact genetic algorithm (ECGA) (Harik (1999)). In the next section, the population-based incremental learning, which is the first developed and most widely applied EDA, is reviewed in detail.

3.2 Population-based incremental learning

As described by Baluja (1994), population-based incremental learning is "a method of combining genetic algorithms and competitive learning for function optimization". Therefore, in this section, we first present genetic algorithms and competitive learning, followed by a review of PBIL algorithm.

3.2.1 Genetic algorithm

Genetic algorithms (GA), initially developed by Holland (1975), are a class of automated optimization algorithms based upon the principles of natural selection and genetic recombination. They are usually applied to solve problems that the solutions can be represented by fixed-length strings with finite alphabet.

Genetic algorithms maintain a group of potential solutions, so called a *population*, which is randomly generated at the initial round. Each population member is referred to as a *chromosome*, which defines a potential solution to an optimization problem. The chromosome is often represented by a string of binary alphabet. In each generation, the *fitness* of each chromosome or

potential solution is evaluated, in other words, it measures how well each solution optimizes the objective function. The higher the fitness, the better the candidate solution.

Three operators of the simple genetic algorithm, namely, reproduction, crossover and mutation, make GA an effective technique to generate favorable results in many practical problems. According to Goldberg (1989), "Reproduction is a process in which individual strings are copied according to their *objective function* values, f (biologists call this function the *fitness function*) ... In natural populations fitness is determined by a creature's ability to survive predators, pestilence, and the other obstacles to adulthood and subsequent reproduction. In our unabashedly artificial setting, the objective function is the final arbiter of the string-creature's life or death". Therefore, the fitness function can be considered as a type of objective function which measures the optimality of a solution or a chromosome in an algorithm. In financial research, the objective function can be, for instance, a measure of trading volume, share price, trading cost, profit or utility that we want to minimize or maximize. Copying strings in accordance with the fitness function represents that the population with high fitness values will have a higher probability of being selected as a "parent". This artificial form of natural selection is a typical "Darwinian survival of the fittest among string creatures."

After reproduction, the chromosome of "child" may be slightly different through the process of crossover. First of all, all the reproduced strings in the mating pool are paired at random with each other to merge the information and characteristics of both parents into the subsequent generation. Goldberg (1989) describes the process that "each pair of strings undergoes crossing over as follows: an integer position k along the string is selected uniformly at random between 1 and the string length less one $[1, l-1]$. Two new strings are created by swapping all characters between positions $k+1$ and l inclusively." The *crossover rate* specifies the percentage of the time that the crossover operator occurs when two population are selected to reproduce. If there is no crossover, the two parents pass their entire gene directly into the next generation.

	Parent A	000010	00001	
One point	Parent B	111110	11110	
crossover	Child A	000010	11110	
	Child B	111110	00001	

	Parent A	000	0100	0001
Two point	Parent B	111	1101	1110
crossover	Child A	000	1101	0001
	Child B	111	0100	1110

	Parent A	00000000001
Uniform	Parent B	11111111110
crossover	Child A	01010100011
	Child B	10101011001

	Parent A	001	01100001
"Cut and	Parent B	111011	11110
splice"	Child A	001	11110
	Child B	111011	01100001

Fig. 3.1: Crossover operators used in simple GAs. Modified based on Appendix A of Baluja (1994)

Various crossover methods have been employed for different data structures. And, the crossover type specifies what information is merged between parents. Figure 3.1 shows samples of four crossover techniques: one-point crossover, two-point crossover, uniform crossover and finally the "cut and splice" approach, see Schaffer and Morishima (1987) and Baluja (1994). *One-point crossover*: Given two parent chromosomes, a single crossover point on both parents' strings is randomly chosen. All data beyond the chosen point is swapped between the two parent chromosomes. *Two-point crossover*: Two random crossover points are selected, and the contents of the parents chromosomes between these points are swapped. In the *uniform crossover* schemes, two parents are combined randomly, and each bit position is swapped with a specific probability of 0.5. On the other hand, in the *"cut and splice"* approach, each parent string has a separate selection of crossover point, which may cause a variation in length of the offspring strings.

Mutation is another genetic operator used to preserve diversity from one generation to the following. As explained by Goldberg (1989), "In the simple GA, *mutation* is the occasional (with small probability) random alteration of the value of a string position. In the binary coding of the black box problem, this simply means changing a 1 to a 0 and vice versa". The *mutation rate* is the probability of changing each arbitrary bit from its original state at each generation. The main intention of mutation is to avoid local minima by preventing the chromosomes from being too similar to each other. It, as a result, generates random changes into the offspring chromosomes.

Since parents with higher fitness are more likely to pass their characteristics on to the child, the average fitness would generally increase over a number of generations. Therefore, the potential solutions in the last generation are expected to be the best solution ever found during the optimization process, and the overall trend gets towards fitter solutions. However, due to random factors involved in reproduction, the children may or may not have higher fitness values than their parents.

Though GAs have been applied successfully in many different problem scenarios, see Cliff (2001), Machado et al. (2002), and Cliff (2003), there are still issues that need to be resolved, such as efficient problem representation and adequate scaling of functions to make sure the good genes pass down to the subsequent generation. Among a fixed number of generations, it is difficult for GAs to return the optimal solution due to the randomized searching and the inability of comparing the very small difference between good and optimal solutions. Moreover, due to the mutation and crossover, some characteristics from good solutions may be destroyed as passing to the subsequent generation. And, these solutions may not be found again.

3.2.2 Competitive learning

Competitive learning (CL), clustering a number of unlabeled points into different groups based on the similarity of points, is often used in the field
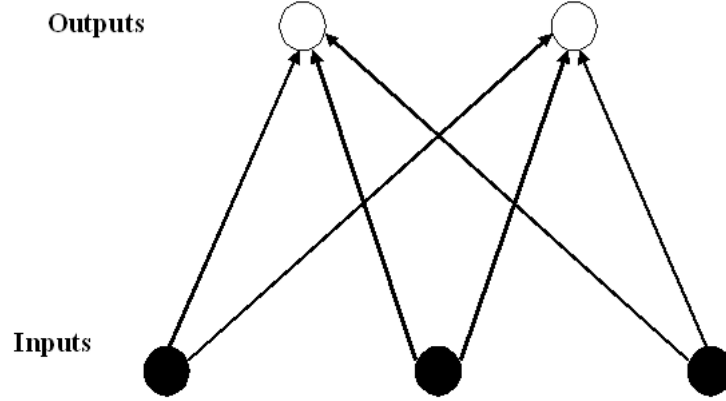
Fig. 3.2: A competitive learning network. Modified based on Figure 1 of Baluja (1994).

of artificial neural networks, see Baluja (1994). Figure 3.2 shows a simple competitive learning network containing three inputs cells and two output neurons. The input neuron, in this network, represents the feature vector for each point. And, the outputs represent the group where the point has been allocated into.

With random initial weights, the activation of the output neuron i is computed during learning with the following formula

$$(3.1) \qquad output_i = \sum_j w_{ij} \times input_j$$

where w_{ij} is the weight of the connection between input unit i and output unit j. The weights of the output unit with the highest activation, i.e. the winning neuron, are modified to represent the current input better:

$$(3.2) \qquad \Delta w_{ij} = \eta \times (input_j - w_{ij})$$

where η is the learning rate. The learning process repeats until the network is stabilized. In this framework, all input units get the same signal, while

only the winning neuron with the highest output signal is allowed to fire, and represent the current input. Through interaction within the network, output neurons learns to represent different groups of the input signal space. After the learning process, which is called *Competitive Learning*, is complete, the weights of each output neuron can be viewed as a prototype vector for a cluster represented by this neuron.

3.2.3 Population-based incremental learning

As a combination of genetic algorithms and competitive learning, population-based incremental learning, which is classified as an estimation of distribution algorithm (EDA), is based on binary searching space to seek the optimum solution using the technique of probability estimation and sampling. In analogy to GAs, the population-based incremental learning (PBIL) algorithm maintains a population of potential solutions evolving over a number of generations. However, the PBIL attempts to create a probability vector specifying the probability of each bit position containing a '1' in a binary string to define the population, see, for example, figure 3.3. As shown in the figure, although population 1 and population 3 having the same probability representation, the solution vectors of the two population are totally different from each other. Thus, this algorithm is capable to maintain diversity in search as the same probability vector could generate distinct population

Instead of transforming each individual into a probability vector used for generating and recombination, the probability vector is moved towards representing the high evaluation vector with a similar manner to the competitive learning process. The probability vector is updated based on the following rule

$$(3.3) \qquad\qquad \pi_t = (1 - \eta)\pi_{t-1} + \eta b_i$$

where π_t denotes the probability of containing a 1 in bit position i, η is the learning rate, and b_i represents the i-th bit value of the best solution in the

Population 1	Population 2	Population 3
0 1 0 1 0	1 1 0 0 1	1 0 1 1 0
1 1 0 0 1	0 1 1 0 0	0 1 0 0 0
0 1 1 0 1	1 1 0 0 0	1 1 1 0 1
1 0 1 0 0	1 1 1 0 0	0 1 0 0 1

Representation	Representation	Representation
0.5 0.75 0.5 0.25 0.5	0.75 1.0 0.5 0.0 0.25	0.5 0.75 0.5 0.25 0.5

Fig. 3.3: The probability representation of 3 populations of 5 bit solution vectors. This is modified from original figures provided by Baluja (1994)

current generation according to the fitness function of the optimization.

At the beginning, information about values of solution vectors is not available, hence, the initial probability of each bit position is considered as 0.5, which means at each bit position, either 1 or 0 is randomly created with equal probability. However, the values of the probability continuously vary towards values representing better evaluation solutions through an evolution process over T generational cycles. At each generation, after sampling a set of n candidate solutions by the current probability vector, each solution is evaluated according to the fitness function, and the best solution among the samples is selected. At the next iteration, the probability at each bit position in the probability vector is pushed towards the best solution with the highest fitness from the previous generation, using formula 3.4, to ensure the best solution during the search is retained. The distance that the probability changed is determined by the learning rate.

To maintain the diversity, mutation is operated in PBIL in an analogous manner as it is applied in GA. Baluja (1994) introduces two ways to perform the mutation before sampling to the next generation. One is to directly mutate the vectors obtained, while the other is to mutate the probability vector, with the following formula

$$(3.4) \qquad \pi_t = (1 - \epsilon)\pi_t + \omega\epsilon$$

where ϵ represents the mutation rate, and ω is either 0 or 1 in Baluja (1994). New samples of the solutions are generated after updating the probability vector. And, the whole process repeats insistently. As it continues, the probabilities change from the initial set up of 0.5 towards either 0 or 1. When the searching loop stops, either the convergence is satisfied which implies that there is no more improvement in the solutions, or the maximum number of iterations is reached, or values in the probability vector are either 0 or 1. The result in the final iteration is therefore considered to be the best solution of the optimization problem.

3.2.4 Further developments of population-based incremental learning

3.2.4.1 Multiple populations

Following the research of Branke et al. (2000), which confirm that using multiple-populations rather than one population works effectively to improve the performance of Evolutionary Algorithms on dynamic optimization problems, Yang and Yao (2003) introduce multiple-populations into population-based incremental learning, and propose two variations of PBIL by employing multiple probability vectors. The first variation is Parallel Population-Based Incremental Learning, which has two parallel probability vectors. In order to make a comparison, one probability vector, applying the same rules as in the original PBIL, is initialized with probability 0.5, while the other probability vector is randomly generated. Both are sampled and updated independently. The sample sizes for both test are equal at the initial iteration, which is half of the total population. Then, the sample sizes would be adjusted accordingly based on the respective performance of each test, i.e., the sample size would be increased if the test performs better than the other one, and the sample size of the other test would be reduced by the same amount.

Another proposed variant is Dual Population-based Incremental Learning, inspired by the nature of complementarity mechanism. Different from the

Parellel PBIL where two probability vectors are updated independently, this approach consists of a pair of dual probability vectors $P = (P_1, ..., P_L)$ and $P' = dual(P) = (1 - P_1, ..., 1 - P_L)$. The two probability vectors are "dual to each other with respect to the central point in the search space". As the principle of dualism, only one probability vector P is updated from learning the best solution, while P' changes automatically. By introducing dualism into PBIL, samples are widely diversified and a higher robustness towards a significant environment changes.

3.2.4.2 Continuous search space

The standard PBIL is proposed for a binary search space or using binary variables. Some recent papers have suggested to extent its application to continuous search space or using a real-coded variant. Two approaches of extensions are reviewed in this section. One approach, proposed by Servet et al. (1997), is to replace the binary strings with an interval $[Low, Up]$. Each gene or variable in the function optimization X contains a "continuous and bounded definition set". Three parameters are specified to determine the value of X, namely, the lower boundary Low, the upper boundary Up, and a probability P interpreted as $P(X > \frac{Low+Up}{2})$, which means the probability of X is uniformly chosen from $(\frac{Low+Up}{2}, Up]$. The probability vector updating rule is same as in the binary version, but with an additional updating of the intervals. Servet et al. (1997) describe the updating rule as following: If the $i - th$ component of P is smaller than 0.1 or larger than 0.9, the Up_i is updated to $\frac{Low_i+Up_i}{2}$, or the $Lower_i$ is updated to $\frac{Low_i+Up_i}{2}$ respectively, and the probability P_i is re-initialized at 0.5. Hence, the new bound of the search interval becomes $(Low, \frac{Low+Up}{2}]$ or $(\frac{Low+Up}{2}, Up]$ correspondingly. This mechanism starts from a most general model and gradually focuses on a small region of the whole search space which contains the best solution in most cases. However, several limitations exist within this framework, such that the sampling is restricted within the bound, and the changes of boundary is irreversible.

Another version of PBIL with continuous search spaces is extended by using Gaussian distributions $N(X, \sigma)$ to model the distribution of the population, and applying the selected solutions to update the distribution, see Sebag and Ducoulombier (1998). The mean of the normal distribution vector, X, is initially set at the center of the search space, and updated iteratively by learning from the two best and the worst solutions at the current generation, see equation 3.5.

$$(3.5) \qquad X^{t+1} = (1 - \alpha) \cdot X^t + \alpha \cdot (X^{best,1} + X^{best,2} - X^{worst})$$

Except for the mean, the standard deviation σ, which measures the diversity of the population, is updated through different ways of incremental learning. Even though the assumptions of a normal distribution can be justified in many circumstances, it may not be accurate in all cases. In this case, the application of this approach is rather limited.

3.2.5 Applications of PBIL and comparison with other algorithms

The population-based incremental learning has been successfully applied in many optimization problems. Baluja (1994) uses PBIL to resolve twelve problems that can be divided into three categories: the NP-complete problems (including Jobshop Scheduling, Traveling Salesman, and Bin Packing problems)[1]; Standard Numerical Optimization Problems (which are usually used to assess the performance of GA in literature), and finally the Strong Local Optima Problems. Based on the empirical findings, Baluja denotes that the population-based incremental learning is capable to solve a wide range of problems. In addition, by comparing the performance of a standard genetic algorithm with a simple PBIL, he indicates that "The simple population-based incremental learner performs better than a standard genetic algorithm in the problems empirically compared in this paper. The results achieved

---

[1] See Baluja (1994) for details.

by PBIL are more accurate and are attained faster than a standard genetic algorithm, both in terms of the number of evaluations performed and the clock speed."

Following the research of Baluja (1994), many studies have been conducted to explore the PBIL evolutionary algorithm. For instance, Southey and Karray (1999) show that a variant of PBIL, which is developed with floating-point representations rather than bit-strings, achieve better results than a standard genetic algorithm in evolving a neural networks-based controller for robotic agents. The studies of Baluja and Caruana (1995) and Larranaga and Lozano (2002) also confirm that the PBIL outperforms on a large number of real-word problems.

A recent study of Yang and Yao (2005) evaluates and compares the performance of a traditional GA and PBILs, which include the basic PBIL, Parallel PBIL, and Dual PBIL. They carry out experimental tests with a set of stationary and dynamic optimization problems. Based on the results, they indicate that the performance of all three PBIL variants have been found to achieve superior performance compared to the standard GA on the stationary problem. However, with the mutation process, the diversity in the population improves the SGA performance on dynamic problems. Besides, although the use of multi-population into PBIL may not improve the performance on stationary problems, it does significantly improve the performance on dynamic problems. The introduction of dualism into PBIL also leads to an enhanced adaptability, and better performance in a dynamic environment.

To examine the performance and adaptability of evolutionary optimization algorithms, Gosling et al. (2004) conduct a comparative study on the performance of a simple genetic algorithm and population-based incremental learning involving the Iterated Prisoners Dilemma (IPD), which is initially studied in Axelrod (1987). In contrast, this experimental result suggests that a GA approach is slightly outperform to a PBIL approach when both framework are stabilized. Although increasing the PBIL population size or reducing GA population size could improve the performance of PBIL com-

paring with GA, employing mutation or a high learning rate seems not useful to close any perceived performance gap.

3.3 Multi-objective optimization

In the previous sections, we have presented a review of evolutionary algorithms for single objective optimization. Although some real-world problems can be solved in a matter of a single objective, quite often it is difficult to measure the quality of the performance based on one objective function since the design may require different features which may be competing to each other. Thus, most of real world optimization problems involve simultaneous optimization of multiple incommensurable and often competing objectives, or measures of performance. To maintain the design quality, all of these objective functions should be considered. For example, financial market participants often value different aspects of an exchange market and will thus likely to use different performance measures, requiring the use of various objective functions to optimize the market structure. These performance measures can include the volatility of prices, market liquidity, bid-ask spread, trading costs, the information efficiency of prices, allocational efficiency of the market, the profitability of trading, etc. Given that all types of market participants will be active in the market, all their concerns have to be taken into account when determining the optimal market structure. Instead of considering only one aspect of performance such as maximizing trading volume in the market, other aspects of performance functions are required to be optimized simultaneously to represent the different interests of various stakeholders.

The notion of applying evolutionary computation techniques to solve multicriteria optimization problems (MOP) dates back to the pioneer work by Rosenberg (1967). The so called evolutionary multiobjective optimization is very important. It has been used by engineers, computer scientists, biologists, and operations researchers alike. Its application in real-world prob-

lems is continuously increasing, which can be found in various fields: finance, product and process design, automobile design, the oil and gas industry, or wherever multiple conflicting objectives need to be optimized for decision making, such as maximizing profit while minimizing the production cost, see Coello Coello et al. (2002).

Issues related to the multi-objective optimization (MO) are discussed in this part. Section 3.3.1 provides some basic concepts and terminology. Traditional approaches to approximate the set of Parato-optimal solutions are listed in section 3.3.2. This is followed by a discussion of various multi-objective evolutionary algorithms in section 3.3.3. Section 3.3.4 presents the population-based incremental learning for multiobjective optimization, including its methodology and empirical application. In section 3.3.5 the convergence of MO is discussed. Finally, section 3.3.6 shows some empirical comparison of various evolutionary algorithms of multiobjective optimization.

3.3.1 Basic concepts and terminology

General MOP. The *multiobjective optimization problem*[2], which is also called multi-criteria, multi-performance, multi-attribute or vector optimization problem, is defined by Steuer (1986) and Sawaragi et al. (1985) as a problem of optimizing two or more conflicting objectives simultaneously with certain constraints. Osyczka (1985) describes it in details that MOP is the process of findings:

> "a vector of decision variables which satisfies constraints and optimizes a vector function whose elements represent the objective functions. These functions form a mathematical description of performance criteria which are usually in conflict with each other. Hence, the term 'optimize' means finding such a solution which

---

[2] Note that the terms "multi-objective" and "multiobjective" are used interchangeably in this thesis

would give the values of all the objective function acceptable to the designer."

As presented by Zitzler (1999) and Coello Coello et al. (2002) , the multiobjective problem contains n parameters or decision variables, a set of k objective functions, and several equality and inequality constrains. Both objective functions and constraints are functions of the decision variables. Without loss of generality, a MOP can be written in the following form:

(3.6)
$$
\begin{aligned}
\text{maximize} \quad & y = f(x) = (f_1(x), f_2(x), ..., f_k(x)) \\
\text{subject to} \quad & g_i(x) \geq 0 \quad i = 1, 2, ..., m \\
& h_i(x) = 0 \quad i = 1, 2, ..., p \\
\text{where} \quad & x = (x_1, x_2, ..., x_n) \in X \\
& y = (y_1, y_2, ..., y_k) \in Y,
\end{aligned}
$$

in which x represents the decision or parameter vector, $g(i)$ and $h(i)$ are constraints, y denotes the objective vector, X is the decision space, and Y is the objective space. In general, these objectives cannot be optimized at the same time due to the trade-off [3] relationship between them. As a consequence, the main task of a multiobjective optimization is to find a set of Pareto-optimal solutions. We explain these concepts in details in the following part.

Decision variables. The *decision variables* are described by Coello Coello et al. (2002) as the "numerical quantities for which values are to be chosen in an optimization problem". In mathematical terms, each numerical quantity is represented by x_j, $j = 1, ..., n$. The vector of decision variables can be

---

[3] The term "trade-off" in this thesis implies that improvement in one objective function cannot be achieved without worsening another one.

written as:

$$(3.7) \qquad X = \begin{bmatrix} x_1 \\ x_2 \\ . \\ . \\ . \\ x_n \end{bmatrix}$$

or

$$(3.8) \qquad X = [x_1, x_2, ..., x_n]^T,$$

where T denotes the transposition of the column vector to the row vector.

Objective functions. To determine how "good" a specific solution is, we need some criteria or measurement to evaluate it. These criteria, specified as computable functions of the decision variables, are referred to as *objective functions*. The term *objective space* or *objective function space* indicates the coordinate space within which vectors resulting from assessing possible solutions are plotted. A number of objective functions are involved in most real-world multi-objective optimization problems. Sometimes, they may be in conflict with each other, in addition, while some objective functions are maximized others may be minimized at the same time. Keeney and Raiffa (1976) reveal a number of desirable characteristics of the objective functions. Tan et al. (2005) summarize them in the following points: "the objective functions should be

1. Complete so that all pertinent aspects of the decision problem are presented.
2. Operational in that they can be used in a meaningful manner.
3. Decomposable if desegregation of objective functions is required or it is desirable.
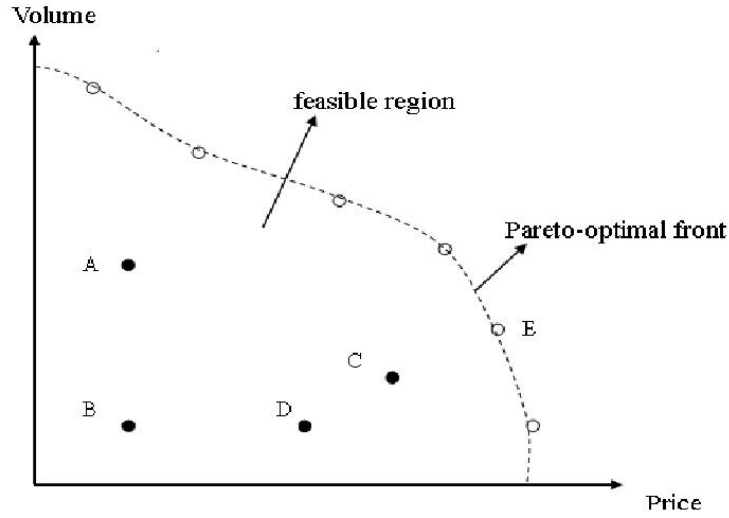
Fig. 3.4: Example of Pareto-optimal solutions in objective space. A modified figure from Zitzler (1999).

4. Nonredundant so that no aspect of the decision problem is considered twice.

5. Minimal such that there is no other set of objective functions capable of representing the problem with a smaller number of elements."

Constrains. In most optimization problems a number of *constraints* or restrictions, which must be fulfilled to consider a feasible solution, are imposed in consideration of certain characteristics of the environment, such as time restrictions, resources limitations, etc. They can be expressed in form of either inequalities or equalities:

(3.9)
$$g_i(x) \leq 0 \quad i = 1, ..., m$$
$$h_j(x) = 0 \quad j = 1, ..., p$$

The number of equality constraints should not exceed the number of decision variables, otherwise, it is overconstrained, as there would be more unknown variables than equations.

Feasible set. The *feasible set* X_f is defined as a set of decision vectors x that satisfy the constraints $g(x)$ and $h(x)$:

$$(3.10) \qquad X_f = \{x \in X \mid g(x) \leq 0 \ \wedge \ h(x) = 0\}$$

For single objective optimization problems, the feasible set is sorted based on the only objective function f: for any two solutions $a, b \in X_f$, either $f(a) \geq f(b)$ or $f(b) \geq f(a)$. Therefore, the goal is simply to find the solution with the maximum (or minimum) value of f. However, for multi-objective optimization, X_f is partially sorted. It is difficult to obtain a single solution which simultaneously optimizes all the objective functions. Here is an example: consider the sale of items, the seller always wants to sell more at the highest price. In this case, two objective functions, price and trading volume, are maximized simultaneously, see Figure 3.4. The solution denoted by point C is better than the one denoted by point D, as the former solution sells more items at a higher price. To compare D and B, despite almost equal trading volume, D outperforms at a much higher price. However, it is difficult to compare A and C, since neither one is superior than the other. Even though the solution represented by A sells more products, it sells at lower price than the one associated with C. Similar to the single objective, the comparison can be expressed mathematically using the relations $\geq$, $>$, and $=$. In this example, it hold that $C > D$, $D > B$, as a result, we can say that $C > B$. As illustrated by Zitzler (1999), for a general situation two objective vectors u and v follow the rules[4] that

$$(3.11) \qquad \begin{aligned} u = v \quad & iff \quad \forall i \in \{1, 2, ..., k\} : u_i = v_i \\ u \geq v \quad & iff \quad \forall i \in \{1, 2, ..., k\} : u_i \geq v_i \\ u > v \quad & iff \quad u \geq v \wedge u \neq v \end{aligned}$$

Instead of achieving a single optimum solution, the aim of the multiobjective optimization is to find good compromises or a set of optimal trade-offs.

---

[4] The relations for a minimization problem $(\leq, <)$ are analogical.

The most commonly adopted notion of "*optimality*" is originally proposed by Francis Ysidro Edgeworth in 1881 Edgeworth (1881), in 1896 it was further generalized by Vilfredo Pareto Pareto (1896). This notion is commonly referred to as *Pareto optimality*. Its related concepts such as *Pareto dominance, Pareto optimality, Pareto optimal set, and Pareto front* have been widely applied and reviewed by researchers, see Zitzler (1999), Van Veldhuizen (1999), and Coello Coello et al. (2002). These terminologies are formally defined as follows (assuming maximization in all objectives).

Pareto dominance. With no available information on the preference of objectives, *Pareto dominance* is applied to evaluate the relative fitness between any two potential solutions in MOPs. This concept together with Pareto optimality have been widely applied. A vector $a = (a_1, ..., a_k)$ is said to dominate vector $b = (b_1, ..., b_k)$ if and only if a is partially more than b. For any two decision vectors a and b in a MOP, there are three possible relations (see Zitzler (1999)):

$$
\begin{aligned}
&a \succ b \ (a \text{ dominates } b) && \text{iff} \quad f(a) > f(b) \\
(3.12) \quad &a \succeq b \ (a \text{ weakly dominates } b) && \text{iff} \quad f(a) \geq f(b) \\
&a \sim b \ (a \text{ is indifferent to } b) && \text{iff} \quad f(a) \not\geq f(b) \wedge f(b) \not\geq f(a)
\end{aligned}
$$

Pareto optimality. A solution $x \in X_f$ is considered to be *Pareto optimal* with respect to a set $A \subseteq X_f$ if and only if there does not exist another $x^\star \in X_f$ such that $f_i(x^\star) \geq f_i(x)$ for all $i = 1, ...k$ and $f_j(x^\star) > f_j(x)$ for at least one j, see Abraham and Goldberg (2005). The "Pareto optimal set" is considered regarding the entire decision space unless otherwise specified. In other words, a solution is Pareto-optimal if there exists no feasible vector which can improve the performance without causing a simultaneous degradation in at least one other objective.

Still referring to the above example and Figure 3.4, we notice that the decision vector represented by C dominates solutions associated with points B and D, and it is indifferent from solution A, yet it is dominated by solution

E. Different from solution A, B, C and D, decision vector of E is not dominated by any other decision vector, i.e., the decision vector is optimal that it cannot be improved without sacrificing another objective. Therefore, all the solutions represented by unfilled circles in Figure 3.4, including point E, are Pareto-optimal, and there are no differences among them.

Pareto optimal set and front. As illustrated by Zitzler (1999), the function $p(A)$ provides the set of Pareto-optimal vectors in A:

$$(3.13) \qquad p(A) = \{a \in A \mid a \text{ is nondominated regarding } A\},$$

with $A \subseteq X_f$. Pareto optimal solutions are those solutions whose objective vector components cannot be all improved at the same time within the decision space. These solutions are also known as *non-inferior*, *admissible*, or *efficient* solutions, while the vectors corresponding to the solutions in the Pareto-optimal set are also referred to as *nondominated*. The entire set formed by all solutions whose associated vectors are nondominated is Pareto optimal set. And, the curve generated by joining the Pareto optimal set under the objective function is called Pareto optimal front, see figure 3.4.

3.3.2   Traditional approaches

Relative few algorithms have been employed for multiobjective optimization compared to the techniques available for single objective optimization. Traditional approaches usually linearly combine the multiple objectives into a single objective using a particular aggregating function to implement the search for optimal solutions, see Coello Coello (1998), Zitzler (1999) for a brief review. In this section, two conventional approaches for MOPs are discussed, including the weighted-sum method and the constraint method.

3.3.2.1 Weighted-sum method

This is possibly the most intuitive and most frequently applied traditional approach for multiobjective optimization problems. Marglin (1967) and Major (1969) were the first to recommend the use of weighting in multiobjective investment problems. In this approach, the MOP is transformed into a single objective optimization problem, by employing a scalar weight to each objective in order to construct a single aggregate objective function:

$$(3.14) \quad \begin{aligned} maximize \quad & y = f(x) = w_1 \cdot f_1(x) + w_2 \cdot f_2(x) + ... + w_k \cdot f_k(x) \\ subject\ to \quad & x \in X_f \end{aligned}$$

where w_i are the weighting coefficients indicating the relative importance of the objectives. The total weight is always assumed to be:

$$(3.15) \qquad\qquad \sum_{i=1}^{k} w_i = 1$$

Clearly, the solution of a MOP solved by equation 3.14 significantly depends on the relative values of the weight specified to each objective. With very little knowledge on how to choose the weights, the designer has to select the most appropriate weights based on his/her intuition. This technique is hence essentially subjective, and any optimal solution obtained is a function of the coefficients applied to combine the objectives. Furthermore, this method cannot identify all the Pareto-optimal solutions. As reviewed by Coello Coello (1998), "This approach is very simple and easy to implement, but it has the disadvantage of missing concave portions of the trade-off curve (in other words, the approach does not generate proper Pareto optimal solutions in the presence of non-convex search spaces), which is a serious drawback in most real-world applications"

3.3.2.2 Normal constraint method

The normal constraint approach[5], proposed by Marglin (1967) to solve multi-objective problems, maximizes the single most preferred or primary objective function, and transforms the remaining $m - 1$ of the m objectives into con-strains:

$$
\begin{aligned}
maximize \quad & y = f(x) = f_h(x) \\
subject\ to \quad & e_i(x) = f_i(x) \geq \epsilon_i \quad (1 \leq i \leq m,\ i \neq h) \\
& x \in X_f
\end{aligned}
\tag{3.16}
$$

The ϵ_i specifies the lower bounds, which are adjusted by the decision maker in order to obtain the multiple nondominated solutions. If the defined lower bounds are not appropriate, there might be no solution to the transformed single optimization problem. To prevent this, a proper range of values for the ϵ_i needs to be chosen. However, the required knowledge may not be available. Another weakness of this approach is that if there are too many objectives, the coding of the MOPs may be difficult or impossible to write. In addition, solutions obtained by this approach are more likely to be weakly non-dominated solutions, which may be unsuitable in certain applications.

3.3.3 Evolutionary multi-objective optimization algorithms

Conventional optimization techniques often have difficulties in solving mul-tiobjective optimization problems, e.g., losing the optimal solutions at the concave portions of a tradeoff curve, see Coello Coello (1996). The weak-ness of traditional methods has motivated researchers to develop alternative methods to find Pareto optimal solutions. Evolutionary algorithms have been recognized to be very efficient and effective for solving sophisticated multi-objective optimization problems where conventional optimization techniques

---

[5] This is also referred to as ϵ-constraint approach,

fail to work well. A multi-objective evolutionary algorithm (MOEA), incorporating the concept of Pareto optimality, evaluates performance of a set of possible solutions at multiple points simultaneously. It allows to obtain the whole set of Pareto optimal solutions in a single simulation run instead of performing a series of separate runs in conventional methods. Evolutionary algorithms have been applied to MOPs for more than two decades. The numerous applications and the rapidly growing interest in the multiobjective evolutionary algorithms have been gaining an increasing attention, see Fonseca and Fleming (1995b), Coello Coello (1998), Zitzler (1999), Tan et al. (2005), and Zhang and Li (2007). We have reviewed evolutionary algorithms for single objective problems in the preceding section. This section briefly reviews some of the most popular evolutionary approaches to multiobjective optimization.

3.3.3.1 Vector Evaluated Genetic Algorithm

Schaffer (1985) is generally considered the first implementation of a multiobjective optimization evolutionary algorithm. Schaffer proposes vector evaluated genetic algorithm (VEGA) by extending the Grefenstette's GENESIS program which is based on a Simple Genetic Algorithm (SGA). The VEGA considers the multiple objectives individually during the evolution. It repeats the selection procedure for each separate objective to obtain a fraction of each succeeding populations to fill up the mating pool[6]. For example, in a MOP with k objectives and a population size of P, k subpopulations each consists of P/k individuals would be generated. Every individual in subpopulation is evaluated, and selected based on one of the k objective functions to fill up a fraction of the mating pool. All the subpopulations are shuffled together to generate a new population, which is then followed by the genetic operations of crossover and mutation. This process is repeated until a predetermined stopping criterion is met.

---

[6] As described by Deb (2008), the mating pool is "an intermediate population (usually created by the selection operator) used for creating new solutions by crossover and mutation operators."

3.3.3.2 Multiobjective Genetic Algorithm

Fonseca and Fleming (1993) propose a variation of Goldberg's Genetic Algorithm, called Multi-objective Genetic Algorithm (MOGA). In this approach, all nondominated individuals are assigned the same smallest rank value of one; while other population members are ranked according to the number of chromosomes by which they are dominated in the current generation. For example, an individual m_i at generation t is dominated by q_i individuals in the objective domain. The current rank of individual m_i is thus assigned by the following rule

$$rank(m_i) = 1 + q_i \tag{3.17}$$

After all population members are ranked, a raw fitness is assigned to each individual by interpolating from the best (rank=1) to the worst (rank $\leq N$) with a certain function (usually linearly). The fitness value is then calculated by averaging and sharing the raw fitness values among population members with the same rank. This is followed by the selection process based on the ranks to form the mating pool using stochastic universal sampling.

3.3.3.3 Niched Pareto Genetic Algorithm

Niched Pareto Genetic Algorithm (NPGA), which combines a tournament selection scheme and the concept of Pareto dominance, is introduced by Horn and Nafpliotis (1993). In this approach, two randomly selected individuals are compared against a subset from the population (usually about 10% of the entire population). If one of them is dominated by the members of the comparison set and the other is non-dominated, then the non-dominated solution wins. If both are either dominated or nondominated, fitness sharing will be utilized to further differentiate between them. Each solution's niche count is the number of individuals that are located within a certain distance

from the evaluated solution in the variable space. The individual with least niche count is then selected for reproduction.

3.3.3.4   Nondominated Sorting Genetic Algorithm

Srinivas and Deb (1994) established Nondominated Sorting Genetic Algorithm (NSGA), in which the only operation varied from simple genetic algorithm is the selection assignment. All the nondominated solutions in the population are first of all classified into one category with a large assigned dummy fitness value. In order to maintain the diversity, sharing is then achieved by performing selection procedure using degraded fitness values computed by dividing the assigned fitness by the niche-count or the number of individuals around it within a distance. After that, the nondominated individuals in the first class are ignored and the second class of nondominated individuals are identified in the same way, but with a smaller assigned dummy fitness value. The process continues until all individuals are classified with a fitness value. The entire population is then reproduced based on the fitness value.

3.3.3.5   Strength Pareto Evolutionary Algorithms

Strength Pareto Evolutionary Algorithm (SPEA), which integrates different MOEAs, is proposed in Zitzler and Thiele (1999). They suggest to maintain an an external set (archive) storing nondominated individuals of each generation. A *strength* value, similar to the ranking value in MOGA, is estimated for each solution in the external set. In SPEA, the fitness of each individual is measured by the strengths of all external non-dominated solutions that dominate it. The effectiveness of this method, as a result, depend on the size of the external nondominated set. If the size is too large, the selection pressure and thus the searching speed might be reduced. Therefore, the proposers apply a clustering technique called average linkage method to ensure the size below a certain level.

3.3.3.6   Multiobjective Evolutionary Algorithm Based on Decomposition

Multiobjective Evolutionary Algorithm Based on Decomposition (MOEA/D), introduced by Zhang and Li (2007), is a recently developed novel evolutionary algorithm for multiobjective optimization. In this approach, a multiobjective optimization problem is decomposed into a finite number of optimization subproblems. These subproblems are then optimized simultaneously by evolving a population of solutions, while each of them is optimized by using information only from its neighboring subproblems. The population, at each generation, consists of the optimal solution found since the beginning of the run of the algorithm for each subproblem. The neighborhood relations are determined by the distances between their aggregation coefficient vectors.

3.3.4   Multiobjective Population-based Incremental Learning

Another evolutionary algorithm, the population-based incremental learning, has been presented in section 3.2.3 for solving single objective problems. The application of this approach, however, has been further extended by Bureerat and Sriworamas (2007) to multi-objective criteria. In this section, we will review the use of population-based incremental learning in multiobjective optimization problems.

3.3.4.1   Methodology

The standard population-based incremental learning, based upon binary searching space, is developed by Baluja (1994) for single objective design. It has been found to produce better results than a standard genetic algorithm in many cases. Subsequently, it has been further extended by Bureerat and Sriworamas (2007) as a multiobjective optimizer, referred to multiobjective

population-based incremental learning (MOPBIL). Instead of a single probability vector used in the simple PBIL, a probability matrix, which comprises more probability vectors with initial values of '0.5', is created in the MOPBIL in order to diversify the population. As explained by Bureerat and Sriworamas (2007), "Each row of the probability matrix is a probability vector that will be used to create a sub-population. Let, N be the number of design solutions in a population, l be the number of probability vectors and n_b be the number of binary bits. The probability matrix, therefore, has the size of $l \times n_b$ where each row of the matrix results in approximately N/l design solutions as one subpopulation." Having the probability matrix, the binary population is then generated based on the matrix.

Similar to SPEA, an empty external Pareto set is also created in MOPBIL to store the old and current nondominated solutions from the population. The non-dominated solutions obtained from the union set of the current population and former external Pareto set based on their objective values are used to replace non-dominated members in the previous external Pareto set. If the size of the Pareto set exceeds the predefined maximum size, the set will be reduced by removing some solutions using the adaptive grid algorithm. As illustrated by Bureerat and Sriworamas (2007), "By using the adaptive grid algorithm, one of the members in the most crowded region is removed from the archive. The crowded regions are updated and the member in the most crowed region is removed iteratively until the number of non-dominated solutions is equal to the size of the archive."

Having revised the external Pareto set, the probability matrix is then updated based on the nondominated solutions in the Pareto set, and a new population is subsequently generated according to the updated probability matrix. The external Pareto set and probability matrix are improved iteratively till the termination criterion is met.

To update the probability matrix, each row vector of probability matrix is

revised using the same rule applied in the original PBIL

$$\pi_i^{new} = (1 - \eta)\pi_i^{old} + \eta b_i \tag{3.18}$$

where η is the learning rate and b_i is the i^{th} bit of the best binary solution. Different from the PBIL for single objective optimization where b_i can be found straightway, for multiobjective optimization, it is difficult to find the best solution. Two updating schemes are presented by Bureerat and Sriworamas (2007) in order to obtain a suitable value for b_i. To begin with the first scheme, $n_0 < N$ binary solutions are randomly chosen from the current Pareto set. The average value of each bit position of the chosen binary solutions is calculated and utilized as b_i to update the probability matrix. The second scheme uses the weighted-sum method, where weighting factors are randomly created with the condition of $\sum w_i = 1$. The weighted-sum

$$f_w(b) = \sum_{i=1}^{m} w_i f_i \tag{3.19}$$

function (3.19) is used to evaluate each binary solution from the union set of the Pareto achieve and the current population. The solution with the minimum value of the weighted-sum function is selected to update the probability matrix. For both updating schemes, the mutation is also executed following the same procedure as in the standard PBIL approach, which is reviewed in section 3.2.3.

3.3.4.2 Application of Multiobjective Population-based Incremental Learning

The multiobjective population-based incremental learning has been described by Kunakote and Bureerat (2008) as "one of the most powerful and robust multiobjective evolutionary algorithms (MOEAs)". However, Applications of multi-objective optimization using the PBIL algorithm are rare thus far as its development is relatively new.

Kanyakam and Bureerat (2007) apply MOPBIL on vibration design problem of a walking tractor handlebar structure. The results show that "the simple but effective strategy of passive vibration control of the structure" has been obtained using the algorithms. Kunakote and Bureerat (2008) employ MOPBIL to solve a multiobjective optimization problem of an unconventional topological design problem for cantilever plate. The optimization results show that "MOPBIL is a powerful tool for multiobjective topology optimisation although its convergence rate is not as good as OCM[7]. With the use of MOPBIL, an unconventional design problem can be dealt with". Another experimental study of Kim et al. (2009) confirms that "solutions from MOPBIL were closer to Pareto optimal front than those from NSGA-II. In real experiments, the proposed MOPBIL algorithm efficiently provided better solutions with respect to the multi-objectives of the path planner."

3.3.5 Convergence of evolutionary multiobjective algorithms

In multiobjective design, no single optimum solution typically exists but in most cases a set of optimum solutions which represents trade-offs between the different objective functions is observed. Therefore, guiding the search towards the true Pareto optimum front is one of the tasks that a evolutionary multiobjective algorithms should fulfill. In recent years, the convergence to the global Pareto-optimal front becomes an interesting topic investigated by some researchers, see, for example, Hanne (1999), Rudolph and Agapie (2000), Marti et al. (2009) and Goel and Stander (2009).

As denoted by Aittokoski and Miettinen (2008), "It seems not to be widely fathomed that often referred EMOs, such as Pareto Archived Evolution Strategy (PAES), Strength Pareto Evolutionary Algorithm (SPEA) or Elitist Non-Dominated Sorting Genetic Algorithm (NSGA-II) are not guaranteed to converge (Laumanns et al. (2002)). Rather, at some point of the solution process

---

[7] OCM represents optimality criteria method.

the populations start to oscillate ... this behavior has been recognized during the last decade, for example, by Deb (2001) (S. 6.2.5), Hanne (1999), Laumanns et al. (2002), Rudolph (1998), Rudolph (1999) and Rudolph and Agapie (2000)"

Goel and Stander (2009) indicate the difficulty of achieving the convergence in multiobjective optimization by demonstrating that "the GA requires tens of thousands of simulations to converge to the global Pareto optimal front". And, Deb (1999) discusses four features presented in a multiobjective design problem that may cause difficulties in converging to the true Pareto-optimal front, which include multi-modality, deception, isolated optimum, and collateral noise (see Deb (1999) for detailed explanation). In addition, Laumanns et al. (2002) argue that the missing operator for elite preservation is the reason that three MOEAs, namely MOGA, NSGA, and NPGA, could not achieve the convergence to the global Pareto-optimal set.

Some other papers, such as Costa and Oliveira (2002), Abido and Bakhashwain (2003), Kazancioglu et al. (2003), Lee et al. (2004), Costa and Oliveira (2004), Mostaghim and Teich (2004), Kanyakam and Bureerat (2007), Kunakote and Bureerat (2008), Ayala and dos Santos Coelho (2008), Wang et al. (2009), and Kim (2010) simply determine a fixed number of generations after which the optimization experiments are stopped and the results treated as if convergence had been achieved. For instance, the stopping criterion of Costa and Oliveira (2002) is to stop the execution after 250 generations. Lee et al. (2004) terminates the algorithm if the number of the generation reaches 300.

3.3.6 Empirical comparison of evolutionary multiobjective optimization algorithms

A variety of evolutionary algorithms have been developed in solving multicriteria optimization problems for the past two decades. And, the question of which method is superior than others is consequently raised. Many researchers evaluate and compare the performance of different MOEAs to

answer the question. For instance, Van Veldhuizen (1999) measure the performance of four MOEAs with six test functions and a set of performance metrics. And, Knowles and Corne (2000) evaluate the performance of NPGA, NSGA, and PAES (Pareto Archived Evolution Strategy, which is an MOEA proposed by the authors).

Zitzler et al. (2000) provide a systematic comparison of eight evolutionary algorithms for MOPs on six test functions. These algorithms include: VEGA, NPGA, NSGA, SPEA, HLGA (Hajela and Lin's weighted-sum based approach), RAND (a random search algorithm), SOEA (a single-objective evolutionary algorithm using weighted-sum aggregation) and FFGA (Fonseca and Fleming's multiobjective EA). Based on the results, a ranking of MOEAs with respect to the distance to the Pareto optimal front is presented in descending order of merit: 1. SPEA. 2. NSGA. 3. VEGA. 4. HLGA. 5. NPGA. 6. FFGA. This is much consistent with the findings of Zitzler and Thiele (1999).

After developing population-based incremental learning for multiobjective optimisation, Bureerat and Sriworamas (2007) compare the performances of two PBIL algorithms using different updating schemes with another four MOEAs, namely the NPGA, NSGA-II, SPEA-II and PAES. Based upon the experimental results, the multiobjective PBIL is said to be "one of the powerful tools for multiobjective optimisation". Two MOPBILs with different updating schemes are found as equally good. In most cases, MOPBIL is comparable to the previously developed MOEAs in terms of convergence rate, while the MOPBIL has superior capability in providing population diversity.

Kanyakam and Bureerat (2007) employ three well-known multiobjective evolutionary optimizers, namely NSGA-II, PAES, and SPEA-II to measure the performance of MOPBIL. They observe that the Pareto optimal front obtained from MOPBIL is superior than the others. And, PBIL is considered as the "best multiobjective optimiser for the proposed design problems." Kim et al. (2009) compare the performance of MOPBIL to that of NSGA-II. The results show that "solutions from MOPBIL were closer to Pareto

optimal front than those from NSGA-II. In real experiments, the proposed MOPBIL algorithm efficiently provided better solutions with respect to the multi-objectives of the path planner." To get the comparative performance of PBIL, SPEA and MPSO (multi-objective particle swarm optimization), a recent paper by Bureerat and Srisomporn (2010) explores three multiobjective design problems. Based on the results, PBIL and SPEA are superior than MPSO. PBIL significantly outperform SPEA in one design problem, while SPEA and PBIL are found to be equally good in the other two design problems.

4. Agent-based computational model of artificial financial market

4.1 Introduction

As we saw in Chapter 2, it is extremely difficult to use traditional methods from behavioral finance and auction theory to analyse the complicated financial market mechanism. To overcome the limitations this research develops a new approach through the use of automatons, a large number of design features, a wide range of performance functions and evolutionary optimization algorithms to study and obtain the best found market structure of financial markets. In addition, we aim to develop a generalized framework which can be used to analyse various market structures with the single model.

We develop an agent-based model in which traders use a very simple trading algorithm that does not assume rational behavior or any other optimizing rule. Such zero-intelligence (ZI) traders have been first introduced in Gode and Sunder (1993) with the explicit aim to investigate the importance of the trading rules for the outcomes of trading. The use of appropriate automatons would allow us to focus on the influence the market structure, i.e., set of trading rules, has on the outcomes without any influence from traders' behavior. Using computer simulations of the trading behavior the market performance can easily be determined for a variety of performance measures.

With traders essentially behaving randomly with minimal restrictions, we are able to investigate a wide range of trading rules, e.g., the tick size, degree

of intervention by market makers, priority rules, and market transparency, commonly found in financial markets and conduct research into the design of a call market to obtain the optimal combination of the trading rules. By investigating a large number of design features, we can build markets that are much closer to reality than is usually possible, enabling us to evaluate the optimal market structures under much more realistic conditions than an analytical model would enable us to do.

In this chapter I give an overview of the simulation framework which was used to conduct the experiments that are reported on in later chapters. In section 4.2, I discuss how the artificial market is constructed with an overview of the traders, trading process and trading rules used in the market. In order to prove the credibility of our model, we then test its ability to generate some well-known stylized facts of asset returns in financial markets in section 4.3.

4.2 Simulation framework

In this research we use a call market as an approximation of continuous double auction market. The most common continuous double auction markets are financial markets which also provide the background to this research project. In this market we batch all orders in each time step, where a time step consists of the submission and revision of orders as well as the batching of orders, the determination of the transaction price and execution of the trades. All the matched orders are executed in a single trade. In the following sections, we explain the four elements that compose our model, which include the traders' behavior, determination of transaction prices, cancelation and revision of orders, and trading rules considered.

4.2.1 The behavior of traders

The traders used in Krause (2006) submit random limit buy or sell orders which have a distribution around the long-term fundamental trend for buy

orders or around the price at which the asset has previously been bought for sell orders. This trading behavior was able to reproduce the fat tails and multi-scaling, but failing for other stylized facts. We intend to develop these traders further by integrating elements of the learning process used by the ZIP-traders from Cliff and Bruten (1997), and including a simple feedback mechanism which in other models often generates additional stylized facts, see e.g. Iori (2002).

Different from Krause (2006) in which all trades are executed in a double auction market, we investigate a market in which a fixed number of N traders trade a single asset in a call market. We use such a market structure as an approximation for the double auction market that can be found in most markets; consequently we do not cancel orders left unexecuted in the order book but carry these over for consideration in the next trade.

At any time each trader is either a buyer or seller of the asset. Whether a trader is a buyer or a seller is determined as follows: if his last transaction was to buy the asset he becomes a seller and if his last transaction was to sell the asset he becomes a buyer. A change from buyer to seller or vice versa for a trader only occurs if this particular trader has no orders remaining in the order book, i.e., his entire order has been executed. In the initialization of the experiments buyers and sellers are determined randomly with equal probability.

A buyer B_i submits bid orders $B_i^t, i = 1, ..., N$ at time t with the bid price randomly taken from a log-normal distribution:

$$(4.1) \qquad \ln B_i^t \sim N \left(\ln V^t + \mu_{buy}, \sigma_{buy}^2 \right),$$

where V^t is the long-term fundamental value in time period t, and here assume to be equal to the initial price P_0[1]. All traders in this market perceive the same fundamental value of the asset V. Keeping the fundamental value

---

[1] We could also introduce a positive long-term trend of the fundamental value without changing the results of our model.

constant does provide a good approximation in the short time frame we consider here and exogenous changes to the fundamental value are of no interest in our investigation given that market participants, being zero-intelligence traders, are assumed not to have any information they can trade on. μ_{buy} denotes the average amount by which the bid price exceeds the fundamental value, and σ^2_{buy} represents the variance of bid prices around the mean. With $\mu_{buy} < 0$ the limit bid price will on average be below the fundamental value, although traders may well submit orders with limit prices above the fundamental value given the random nature of the limit price in our model. We might interpret this either as uncertainty about the fundamental value to which traders pay limited attention, different opinions about the true fundamental value or the fact that many traders will ignore the fundamental value to a large degree in their decision-making. While experiments have shown that the exact specification of the decision-making process is not affecting results, we require a minimal amount of information which traders use as a common anchor for their decision; this is necessary to avoid the limit prices and thereby transaction prices to evolve such that an infinitely large bubble emerges. This constraint on the behavior of traders thus implicitly acts as a budget constraint as too large limit prices are not permitted thereby limiting the amount that can be invested. Similarly, too small limit prices will not be observed, acting as a minimum size requirement for entering the market.

If we denote by $\widehat{P}_i^{t-1}$ the price at which a trader bought the asset the last time, the ask price quoted by seller S_i at time t is then randomly chosen from the following distribution

$$(4.2) \qquad \ln S_i^t \sim N\left(\ln \widehat{P}_i^{t-1} + \mu_{sell}, \sigma^2_{sell}\right),$$

in which μ_{sell} denotes the average amount by which the ask price exceeds the price previously paid by the trader, and σ^2_{sell} represents the variance of ask prices. A trader will only be able to sell those shares he actually holds, i.e. we do not allow for any short sales, thereby acting implicitly as a budget constraint on the behavior of traders.

The order size for a sell order is always equal to the number of shares held, i.e. a seller must sell all the securities he has in the next round of trading. Therefore,

$$(4.3) \qquad \qquad \ln Q(sell)_i^t = Inv(sell)_i^t,$$

where $Inv(sell)_i^t$ represents the remaining number of shares held by seller i. However, the order size for buy orders $Q(buy)_i^t$ is a random variable with

$$(4.4) \qquad \qquad \ln Q(buy)_i^t \sim iidN\left(\mu_{size}, \sigma_{size}^2\right),$$

where μ_{size} denotes the average of the order size, and σ_{size}^2 is the variance of the order size.

4.2.2   Determination of transaction prices

Following the price formation approach applied in Gode and Sunder (1993) and Cliff and Bruten (1997), the transaction price is determined where the demand and supply curves intersect, i.e., the price at the maximal trading volume is chosen as the transaction price. In this market, limit orders with the highest bid prices are first traded and cleared in the market; oppositely, the cheapest sell orders are traded with priority. If we find that there are multiple prices at which the trading volume shows the same maximal value, we employ trading rules to determine which of the prices will be chosen. As shown in Figure 4.1, when demand and supply curve intersect at the quantity Q, the possible trading prices are P1 and P2. Moreover, any imbalances between buy and sell orders at the transaction price will lead to the need for rationing, see figure 4.2, where D1 is different from S1; how this rationing of buy or sell orders is conducted will depend on the trading rules as outlined below.

Fig. 4.1: Demand and supply schedules for the case of multiple prices. The downward sloping curve represents the demand function and the upward sloping curve represents the supply function. This is redrawn and modified from Cliff and Bruten (1997) Figure 3.
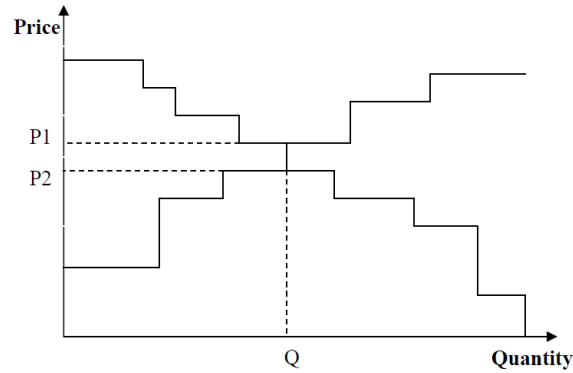

Fig. 4.2: Demand and supply schedules for the case of order imbalance. The downward sloping curve represents the demand function, and the upward sloping curve represents the supply function. This is redrawn and modified from Cliff and Bruten (1997) Figure 3.
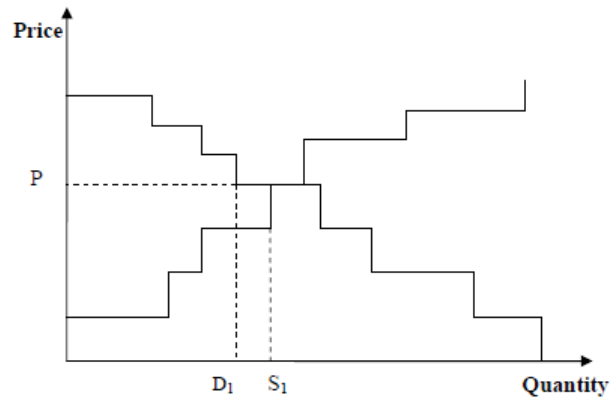
4.2.3   Cancelation and revision of orders

After the completion of each transaction, all the trading information is recorded. The executed orders are canceled from the order book, which contains the traders' ID number, whether they are buying or selling, their limit price, order size, order submission time and length until the order is to be revised. Before the next trading, new limit orders taken from the following distributions are submitted to replace the canceled orders :

$$
(4.5) \qquad
\begin{aligned}
\ln B_i^t &\sim N\left(\ln V^t + \mu_{buy}, \sigma_{buy}^2\right), \\
\ln S_i^t &\sim N\left(\ln P_i^{t-1} + \mu_{sell}, \sigma_{sell}^2\right),
\end{aligned}
$$

where P_i^t denotes the market price at time t.

An order remains in the order book until it is filled or canceled; for partially filled orders the remainder of the order remains in the order book. An order not filled after T_i^t time steps is canceled, where

$$
(4.6) \qquad \ln T_i^t \sim iidN\left(\tau, \sigma_\tau^2\right),
$$

in which τ is the average time of order remains in the order book, and σ_τ^2 denotes the variance of this time. There are at least two reasons for the re-placement. Traders after a period of time may change their initial decision if their orders have not been cleared. The revision of orders could help them to reevaluate and revise their submitted orders. In addition, resubmission could improve the liquidity by traders who need to trade urgently and re-assess their orders to meet their demands. After the cancelation and replacement, the new order book would consist of the remaining orders together with the resubmitted new orders.

4.2.4 Trading rules considered

4.2.4.1 Tick size

In the market we are able to vary a wide range of trading rules. We will firstly investigate different *tick sizes*, i.e. minimum differences between prices at which orders can be submitted. The tick size has several impacts during trading. As it represents the cost of getting inside other competitors' quote, the tick size affects the motivation of submitting limit orders and bid-ask spreads. The empirical literature has been provided in section 2.1.4.1. In order to make limit prices to comply to the tick size, we will lower any limit price of buy orders as determined in (4.1) and (4.5) to the next permissible price and similarly raise the limit price of sell orders determined by (4.2) and (4.5) to the next permissible price.

4.2.4.2 Priority rules

Secondly, different *priority rules* are employed to determine the rationing of orders in the case of an imbalance between buy and sell orders at the transaction price. The enforcement of priority rules, as the primary difference between market structures, is another important design feature of trading systems. We use in particular *time priority*, which is the most commonly used rule. It adheres to the principle of first-come first-served, and ensures that orders submitted earlier will be filled first; *reverse time priority* in which orders submitted later will receive priority to be filled; another frequently used rule to promote traders to place larger orders is *size priority* in which larger orders receive priority; *random selection* in which the orders to be filled are selected randomly and with *pro-rata selection* in which all orders get filled partially to the same fraction.

4.2.4.3 Multiple prices

Thirdly, for the case of *multiple prices* at which the trading volume is maximal we determine the transaction price to be either the price closest to the previous price, the price furthest from the previous price, the highest price, the lowest price, the price with the minimum order imbalance (the absolute value of the difference between the volume of buy and sell orders at the transaction price), the price with maximum order imbalance or a randomly selected price.

4.2.4.4 Market transparency

Fourthly, we also consider *market transparency.* In a transparent market, traders are able to have access to information on the order book and react to any orders submitted by other traders. In order to replicate this aspect of the market we assume that a fraction γ of the traders has access to the order book and can observe the potential transaction price as well as the ensuing order imbalance if the trades were to happen instantly. They use this information to revise their own order size according to the size of the order imbalance (the difference between the size of buy and sell orders at the transaction price) δ for a buy and sell order, respectively:

$$
\begin{aligned}
\widehat{Q}_i^t &= Q_i^t - \alpha\delta, \\
\widehat{Q}_i^t &= Q_i^t + \alpha\delta,
\end{aligned}
\tag{4.7}
$$

where α represents the fraction of the order size revised, Q_i^t is the order size before revision, and $\widehat{Q}_i^t$ is the order size after revision. For instance, if there are excessive buy orders at the transaction price, i.e., $\delta > 0$, an informed buyer will decrease his order size to reduce the order imbalance while an informed seller will increase the order size to absorb the unfilled orders. The revised size is then used to determine the transaction price.

4.2.4.5 Market making

As a final aspect we consider the *intervention of a market maker* into the trading process. A market maker would intervene or influence the prices such that he is prepared to trade a fraction θ of the order imbalance at any time in the market with the existence of imbalance between demand and supply at the transaction price by submitting an offsetting order with price

$$(4.8) \qquad \widehat{\widehat{P}}^t = P^t - \lambda I^t,$$

where I^t denotes the inventory of the market maker, i.e., the number of shares held by him, λ is a coefficient representing the price adjustment of market maker. Holding a volume of shares as inventory, the I^t is a positive number; in contrast, the I^t is a negative value representing a purchasing position. For instance, if the market maker is holding a certain amount of securities, then $I^t > 0$ and the price quoted by market maker can be rewritten as

$$(4.9) \qquad \widehat{\widehat{P}}^t = P^t - \lambda |I^t|,$$

where $|I^t|$ is the absolute value of I^t. In this case, the market maker is willing to set a price lower than market price to stimulate purchasing activity in order to reduce the inventory. On the other hand, if the market maker needs to buy securities to fill up the inventory, then $I^t < 0$ and the price set by market maker would be

$$(4.10) \qquad \widehat{\widehat{P}}^t = P^t + \lambda |I^t|.$$

In this case, the quote by market maker is higher than the market price in order to attract more sellers to enter the market and add additional inventory. To sum up, increases in inventory lead to drop of prices, while inventory reductions result in increase of prices. Such a linear relationship between the price and inventory has been established in the inventory-based models of market-making, see section 2.1.5.2 for an overview.

In line with the behavior of the other traders, we do not assume that the market maker is employing a sophisticated optimization procedure but merely applies a different behavioral rule than other traders. The linear adjustment rule from equation (4.8) as proposed for pure dealer markets does not allow for any strategic interactions of market makers with limit order traders, nor does the price setting ensure that the order imbalance is actually reduced by a fraction θ. Furthermore, we assume that the market maker only submits either a bid or an ask price, depending on which side the order imbalance occurs. In that sense the market maker we introduce here is not directly comparable to those in pure dealer markets but more like those of the specialists of the NYSE who only intervene actively by taking positions in cases of market imbalances, e.g. a significantly larger amount of buy than sell orders (or vice versa) within a given period of time, a comparable situation to the order imbalances in our model.

4.3   Comparison of market structures

The stylized facts of financial assets, including volatility clustering, reversion to the mean, multi-scaling, and fat-tail distribution of returns, which are empirical findings that are robust across a wide range of instruments, markets and time periods, are so widespread that they have become an important benchmark to test any models of trading behavior against their ability to generate these properties. When developing a model of trading which we use to evaluate market structures in this research, it is thus reasonable to investigate its ability to generate the stylized facts through the traders' behavior. Therefore, in this section, we analyze three stylized facts generated by our model, including the long memory of return volatility; fat tails of the return, and multi-scaling of returns. Log return of the asset prices is used as the "return" or the basis for obtaining the stylized facts. For volatility clustering, we suggest to calculate the autocorrelation of absolute returns instead of squared returns. For distribution of returns, we investigate the cumulative

distribution of absolute logarithmic returns on a log-log plot. Finally, we applied the method used in Krause (2006) to analyze the fact of multi-scaling. With various trading rules applied, we can then compare the properties of the resulting asset prices under different market structures.

4.3.1 Parameter constellations considered

We consider a market with 100 traders, which consists of 50 buyers and 50 sellers in the first round. The order book contains the traders' ID number, whether they are buying or selling, their limit price, order size, order submission time and length until the order is to be revised. The initial order book is constructed randomly by determining 50 buy orders with a randomly selected limit price and order quantity using $V^t = 100$; similarly we determine 50 sell orders in the same manner with $\hat{P}_i^{t-1} = 100$. The other parameters chosen are described in table 4.1. We assume that the trading price equals the previous price if there is no trading.

Each simulation is run for 51,000 time steps, where the first 1,000 data is eliminated from the investigation to avoid the influence arising from the initial conditions. To evaluate the effect of various parameter values and priority rules, we analyze the time series with different trading rules. For each simulation only one parameter value or trading rule is changed, while the others remain at the default value.

4.3.2 Price and trading volume

We first present some features of price and volume data obtained from a typical simulation, in which all the trading rules are randomly selected. It is clearly shown in Figure 4.3 that in the simulated market, asset prices fluctuate around the long-term fundamental value of 100. We observe that both high and low prices are temporary and that prices tend to move back towards the average over time, which is consistent with the theory of mean

Parameter	Description	Value
μ_{buy}	Mean of bid exceeding fundamental value	-0.1
μ_{sell}	Mean of ask exceeding fundamental value	0.1
σ_{buy}	Standard deviation of bid	0.3
σ_{sell}	Standard deviation of ask	0.3
τ	Mean time of order in order book	1+ln100
σ_{τ}	Standard deviation of time in order book	1
μ_{size}	Mean of order sizes	1
σ_{size}	Standard deviation of order sizes	1
t	Tick size	0*, 0.1, 0.2, 0.5
α	Fraction of order size revised	0*, 0.02, 0.2, 0.5
λ	Price adjustment of market maker	0*, 0.1, 0.2, 0.5
γ	Fraction of informed traders	0*, 0.2, 0.5, 0.99
θ	Imbalance traded by market maker	0*, 0.2, 0.5, 0.99

* denotes the default value for the experiments.

Tab. 4.1: Parameter values considered in the experiments of comparison of market structures

reversion. In addition, we find that in some regions, prices tend to move slowly and steadily, while in other periods, the amplitude of price fluctuation becomes much higher. It indicates that small price variations are more likely to be followed by small price variations, while large price variations tend to cluster together, which proves the volatility clustering to a certain extent. We also notice that during certain time periods prices remain stable, i.e., the line is flat, which indicates that no trading occurs. This may happen if there are no agreed prices matching between buy orders and sell orders. However, it does not last long as unfilled old orders would be revised to meet the current demand of traders, which may accelerate a matching between the bids and asks.

For trading volume, it is hard to observe a clear trend in Figure 4.4; it seems to move continuously without any pattern. However, compare with the price series, we can see that if the price change, or the absolute value of price increment is large, the market will experience a high trading volume. In accordance with the theory, higher absolute value of price change would increase the volatility, which has a positive correlation with trading volume.
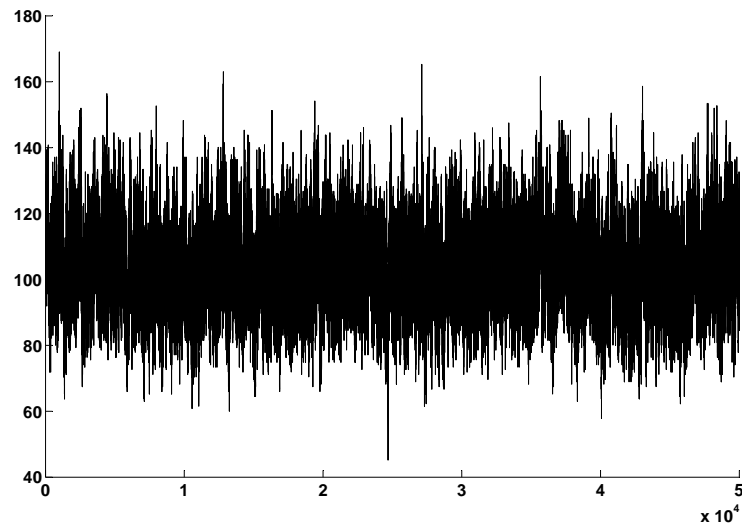
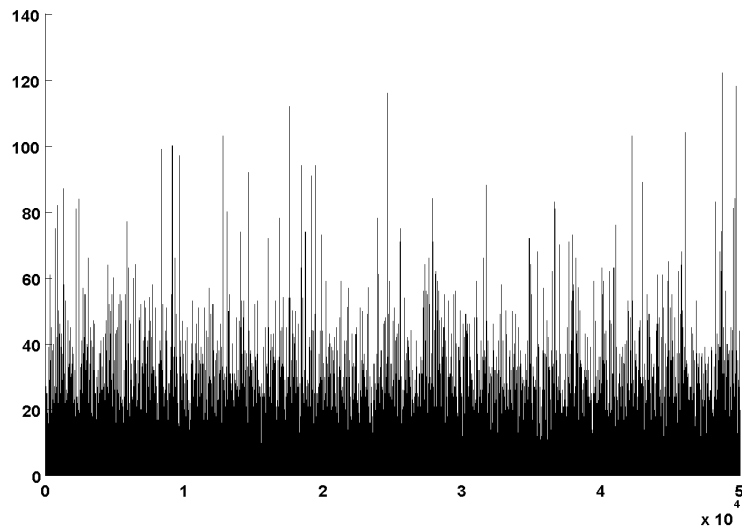**Fig. 4.3:** Price series obtained from a typical simulation

Fig. 4.4: Volume series obtained from a typical simulation

In addition, the trading volume is zero when there is no matching between bids and asks.

4.3.3 Volatility clustering

The autocorrelation function of absolute returns decays only slowly, thus providing evidence for the existence of volatility clustering. The strength of the decay can be measured through the autocorrelation function γ:

$$(4.11) \qquad\qquad \gamma \propto \Delta t^b,$$

where Δt is the lag and b denotes the exponent determining the decay which is empirically smaller and equally to 0.5, (Cont et al. (1997), Breidt et al. (1998)). Based on our results, for most trading rules the exponent takes a value of approximately 0.5 for the investigated parameter settings. However, when we use certain rules to determine the transaction price in cases where multiple prices are feasible, this picture can change. If we choose the price at which the order imbalance is minimal, the exponent will be significantly higher and sensitive to the other parameters. On the other hand, if the price is randomly selected the resulting exponent will be significantly lower, in some instances the obtained exponents are lower than 0.3. This is consistent with **?**, who indicates that "While GARCH models give rise to exponential decay in autocorrelations of absolute or squared returns, the empirical autocorrelations are similar to a power law...with an exponent $\beta \leq 0.5$, which suggests the presence of 'long-range' dependence in amplitudes of returns...".Figure 4.5 illustrates the size of this exponent for representative trading rules. We find that the results are largely stable for a wide range of trading rules and only very few constellations - usually trading rules not commonly found in real markets - give rise to different outcomes.
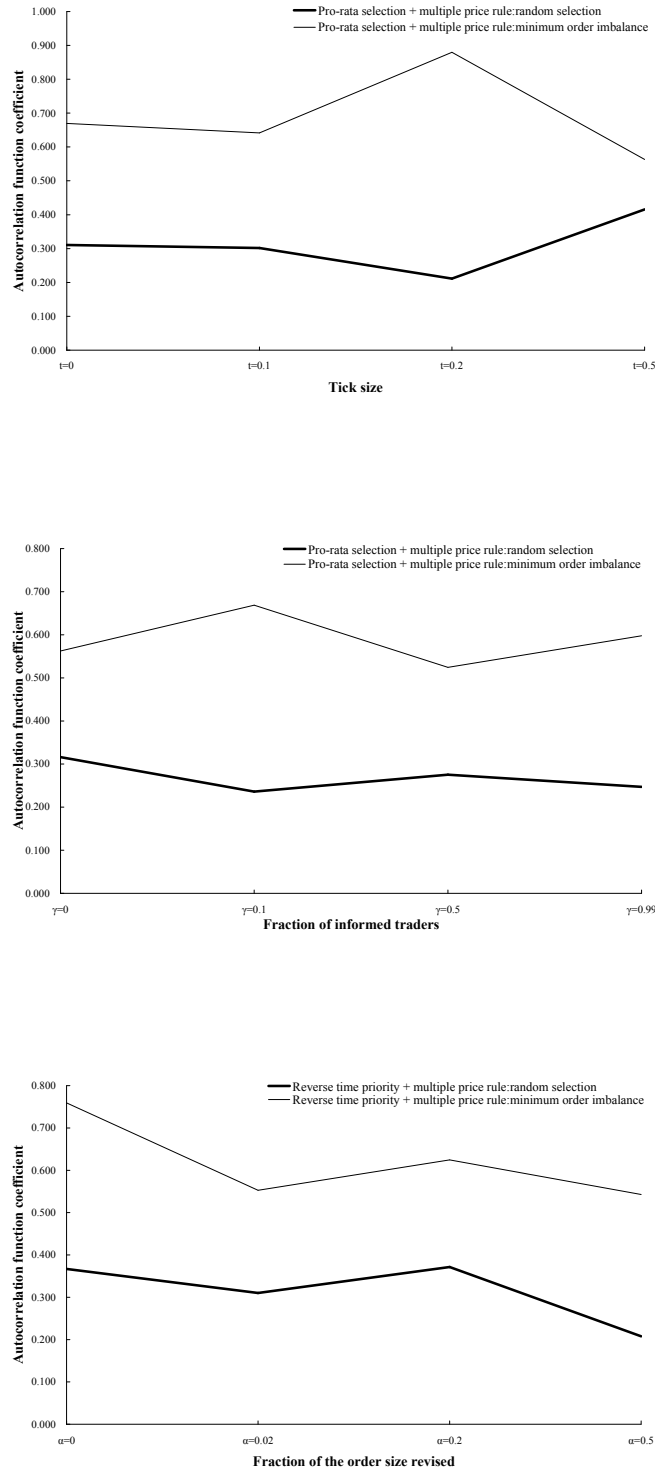
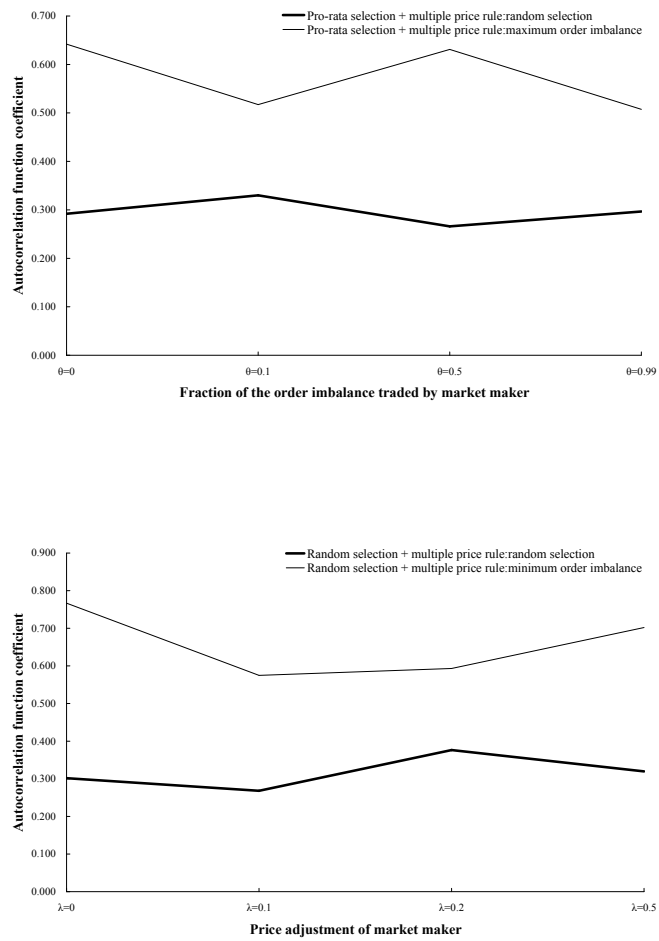**Fig. 4.5:** Volatility clustering coefficient

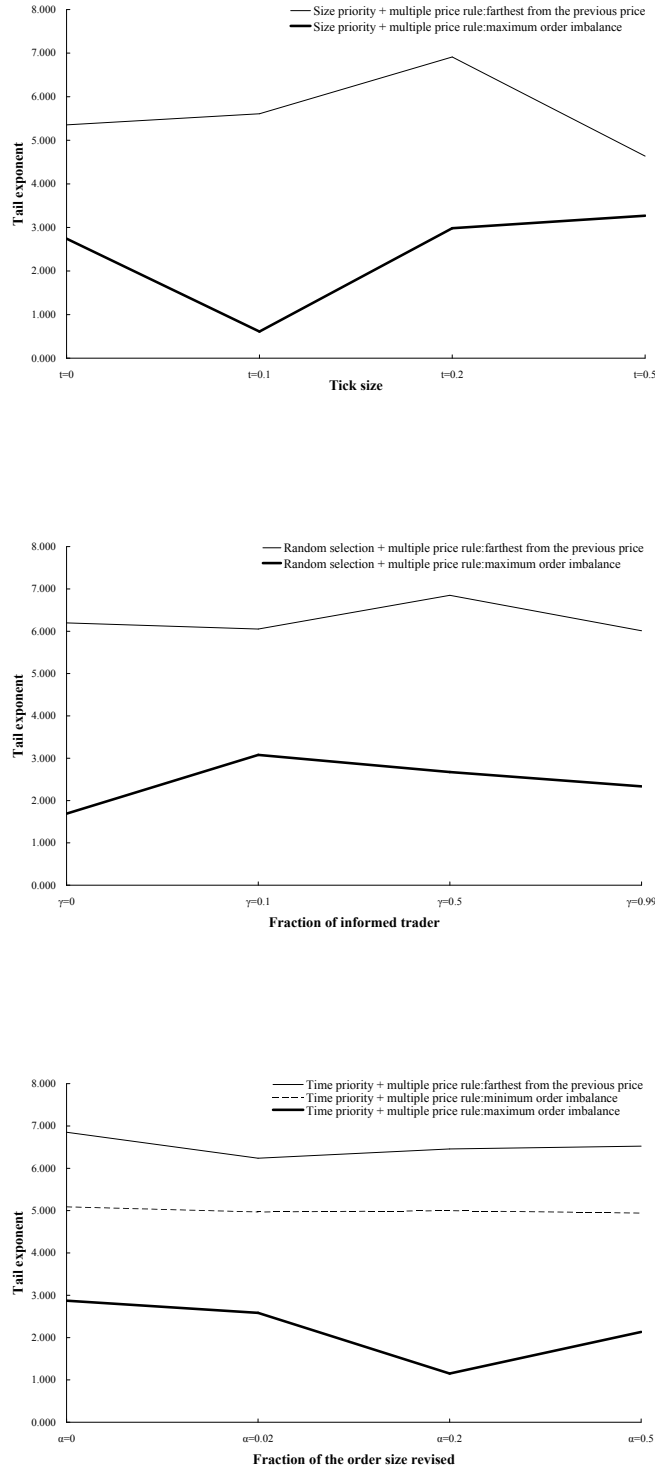**Fig. 4.5:** Volatility clustering coefficient (ctd.)
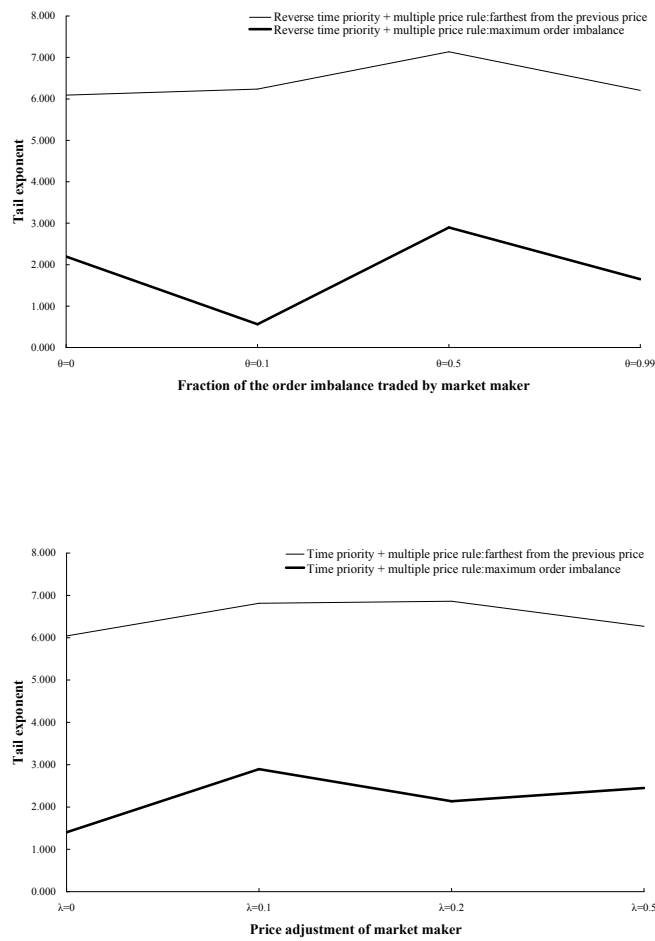
Fig. 4.6: Tail exponents of returns

Fig. 4.6: Tail exponents of returns (ctd.)

4.3.4 Fat tails of the return distribution

One property nearly all financial time series exhibit is fat tails of returns r_t, which manifest themselves in a power law distribution at the tails, showing a tail exponent of $\alpha \approx 3$:

$$1 - CDF \propto |r_t|^{-\alpha}, \tag{4.12}$$

where CDF denotes the distribution function of returns. We find that for most trading rules the tail exponent is approximately 3, and only when the multiple price rule states that the price is determined either randomly or where the minimum order imbalance is, do we find even fatter tails with a tail exponent closer to 1. Figure 4.6 provides representative examples of these findings.

4.3.5 Multiscaling

Another property returns commonly exhibit is multi-scaling. We estimate

$$E\left[|r_t\left(\Delta t\right)|^q\right] \propto \Delta t^{\varphi(q)}, \tag{4.13}$$

where $r_t\left(\Delta t\right)$ denotes the return over Δt time periods and $\varphi(q)$ is a function of the moment q. In the absence of multi-scaling we find $\varphi(q) = qH$, with H being the Hurst coefficient, where $H = \frac{1}{2}$ for normally distributed variables. Multi-scaling emerges for a non-linear function $\varphi(q)$. To investigate the non-linearity of $\varphi(q)$ we run the following regression:
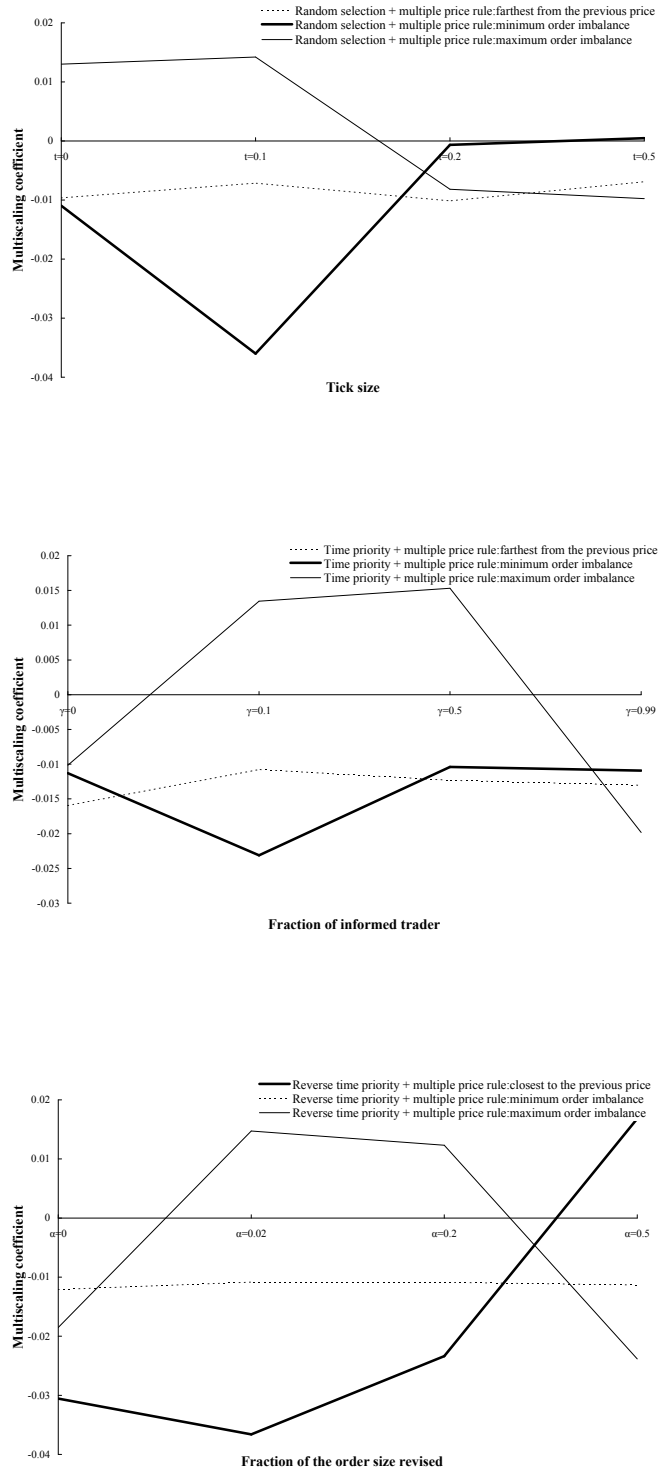
$$\varphi(q) = a + bq + cq^2. \tag{4.14}$$

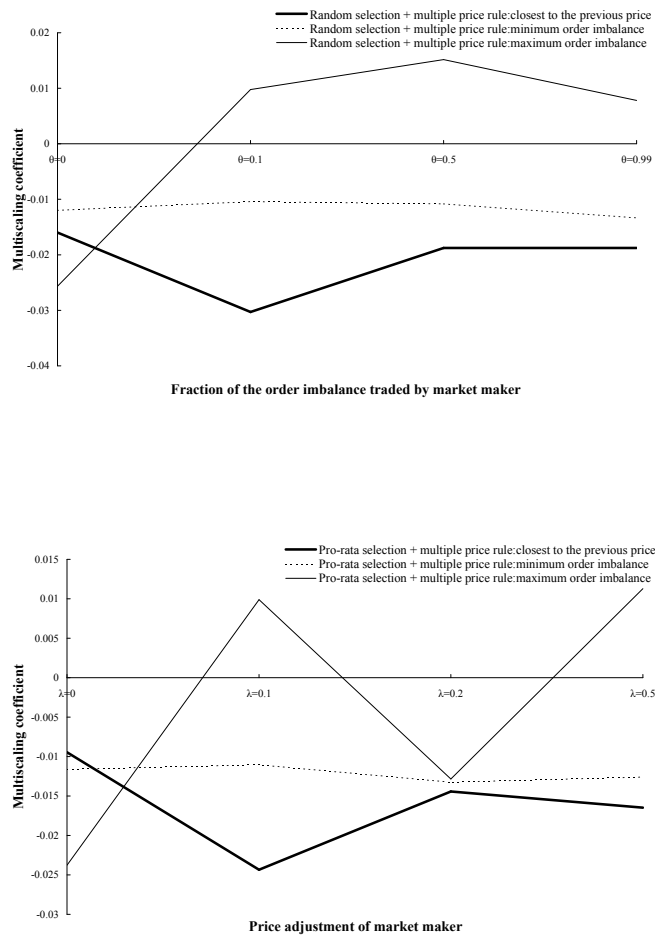**Fig. 4.7:** Multi-scaling of returns

Fig. 4.7: Multi-scaling of returns

If we find that $c \neq 0$, we can conclude that $\varphi(q)$ is non-linear and thus the returns exhibit multi-scaling. We find a very mixed picture regarding multi-scaling, for some trading rules evidence of multi-scaling is very significant while for others it seems to be absent. For the multiple price rule choosing the price with the maximal order imbalance or a random price will produce a convex function $\varphi(q)$. When using priority rules that allocate orders pro rata we tend to find a convex function, while a range of other trading rules do not produce multi-scaling. Figure 4.7 illustrates typical examples of these findings.

4.3.6 Evaluation of computer experiments

Overall we observe that the market mechanism we developed in this paper does reproduce a number of statistical properties that are found in real asset markets, which provides some evidence on the credibility of our model. These properties are found to be robust over a wide range of trading rules and parameter settings. We find that in particular changes in the multiple price rule towards determining the transaction price randomly or choosing the price with the lowest order imbalance seem to change the observed properties of the times series.

We thus can infer that the auction mechanism generates realistic time series properties which are not arising from the complex behavior of agents trading the assets but are rather the result of the market structure itself. It has been found empirically that the properties generated are very stable across asset classes, trading arrangements and time; the robustness of the results in our model with respect to different trading rules are well in line with this result. It suggests that the empirically observed properties are all generated by the auction mechanism and the detailed trading rules employed are only of limited relevance for the observed outcomes.

5. Optimization of market structures with single objective function

Having an appropriate model of simulating financial markets, it is now capable to carry out the investigation of the optimal market structure, i.e. the set of trading rules that provide the best performance. The methodology used to optimize the market structure is a computer experiment in which trading is simulated over a given number of time periods with a given market structure. The optimal trading rules are found evolutionary by population-based incremental learning. As a combination of genetic algorithms and competitive learning it makes PBIL a successful and efficient search mechanism employed in such complex optimization problems as the one presented in this research. Furthermore, in order to consider the different interests of various market participants we employ a number of performance measures as our performance functions, including the trading volume, trading price, return volatility, and bid-ask spread. Using different measures of market performance as objective functions to evaluate market structures we can then analyze the best market structure for various market participants.

In this chapter, section 5.1 describes the optimization framework used for single objective optimization. This is followed by the evaluation of experimental results of simulations with four different objective functions in section 5.2. Finally, section 5.3 concludes the findings of the best-found market structure with single objective function.

5.1 Population-based incremental learning for single objective optimization

In analogy to GAs, the population-based incremental learning algorithm maintains a population of potential solutions evolving over a number of generations. To define a population, PBIL attempts to create a probability vector, measuring the probability of each bit position having a "1" in a binary solution string. Instead of transforming each individual into a probability vector used for generating populations and recombination, the probability vector is moved towards the vector that shows the best performance in a similar manner to a competitive learning process. The probability vector π_t is updated based on the following rule

$$(5.1) \qquad \pi_t = (1 - \eta)\pi_{t-1}^* + \eta\widehat{v}$$

where π_{t-1}^* denotes the probability of containing a 1 in each bit position that was used in the previous generation, and $\widehat{v}$ represents the best solution in the previous generation, selected according to the fitness function of the optimization and η the learning rate. This algorithm is capable to maintain diversity in search as the same probability vector could generate distinct populations.

To employ the PBIL algorithm, we first create a probability vector specifying the probability of each bit having the value one. All initial values for the PBIL process are determined randomly, i.e. the initial probability of each bit is 0.5. With the probability vector thus determined, a population will be produced from this vector according to those probabilities. In each generation we determine the fitness of each potential solution and the best solution in the current generation is selected and used to update the probability vector for the subsequent generation using equation 5.1.

In each time step we determine 100 different parameter constellations using π_t and then determine the best performing parameter constellation from these

100 different market simulations, giving $\widehat{v}$.

Each trading rule is coded into a vector v, where the precision of the continuous variables α, λ, γ, θ is such that each variable is divided into 17 bits each, the tick size t into 20 bits, and the discrete variables (priority rules, multiple prices) are coded such that all rules are covered.

We let this process continue until no changes to $\widehat{v}$ are found, thus the probabilities in π_t are either 1 or 0, giving identical results for all 100 parameter constellations. Therefore, the resulting market structure established by the last remaining set in the process is then the best-found market structure.

Each simulation is run for 2,000 time steps, where the first 1,000 data points are eliminated from our analysis to avoid the influence arising from the initial conditions. The single-objective PBIL optimization is conducted using a population size of 100 over 200 generations with a learning rate of 0.2. We repeat each optimization experiment for 20 times with the same parameter constellation to observe the best-found set of trading rules.

We now investigate the optimal market structure by applying the above single objective PBIL optimization algorithm with four different performance functions. The first objective function is to maximize trading volume, as the higher trading volume, the more liquidity in the market, and orders can be more quickly traded. We also use bid-ask spread as a fitness function. As a measure of the liquidity of the market and of the size of the transaction cost, bid-ask spread is minimised to obtain the optimal market structure. Another objective function is to maximise trading price. For companies going Initial Public Offering, they all wish to be listed in a market where prices keep rising. The last performance measurement we use is return volatility, which is minimised in order to find the best market structure with less risk and uncertainty.

To obtain the optimal market structure with single fitness function, the market and initial order book are constructed based on the same rules and parameter values in section 4.3.1, except that the trading rules are selected

randomly from a specified range rather than choosing from a pool of four different values. The parameters chosen are described in table 5.1 and the initial price P_0 set at 100.

Parameter	Description	Value
μ_{buy}	Mean of bid exceeding fundamental value	-0.01
μ_{sell}	Mean of ask exceeding fundamental value	0.01
σ_{buy}	Standard deviation of bid	0.3
σ_{sell}	Standard deviation of ask	0.3
τ	Mean time of order in order book	1+ln100
σ_τ	Standard deviation of time in order book	1
μ_{size}	Mean of order sizes	1
σ_{size}	Standard deviation of order sizes	1
μ	Trend of the fundamental value	0
t	Tick size	[0,10]*
α	Fraction of order size revised	[0,1]*
λ	Price adjustment of market maker	[0,1]*
γ	Fraction of informed traders	[0,1]*
θ	Imbalance traded by market maker	[0,1]*

* denotes the parameter range used in the optimization.

Tab. 5.1: Parameter values considered in the computer experiments of single-objective optimization

5.2 Evaluation of computer experiments

The results achieved by PBIL are converged fast, normally achieved within 120 generations for all simulations. The convergence is determined based on the criterion that all the probability vector values go to either 1 or 0, which mean there will be no further changes in the results. The experimental results with the objective function of maximizing trading volume are discussed in the following. Results for the remaining experiments with the other objective functions are discussed in Appendix A.

5.2.1   Maximization of trading volume

We first investigate the optimal market structure with the objective function of maximizing trading volume. Trading volume, as an indication of market liquidity, represents the total number of shares or assets traded in a market for a given period of time. Market participants have a preference of markets with higher volume, since the higher the volume, the easier it is to get in and out of a position quickly; the more liquidity is present, and the more competitive the market is. Table 5.2 summarizes the best-found trading rules, as well as the maximal trading volume from the 20 simulation experiments. Due to the stochastic nature of the trading process, even if we use the same combination of trading rules, the resulting trading volume is still different in each run. Thus, to make them more comparable between results from 20 simulations, we use the maximum of the trading volumes that are generated after the results of desirable trading rules converge. The table is sorted by the first column, the trading volume, from smallest to largest. The second column contains "Fraction specialist", which represents the fraction of imbalance traded by market maker. This is denoted as Theta in section 4.2.4.5 on page 117. "Fraction revised" in third column refers to the fraction of order size revised by informed trader. This is denoted as Alpha in section 4.2.4.4. "Price adjustment" in forth column is the coefficient used to get the market maker's quotes, see equation 4.8 in section 4.2.4.5. "Fraction informed" in fifth column denotes the fraction of informed traders in the market, see section 4.2.4.4 for detailed explanation. The sixth column contains the best-found tick size, while the last two column present the best-found multi-price rule and priority rule. All the symbols used here are shown in table 5.1.

After obtaining the best-found market structure, we present the results, the observations and the conclusions of the experiment in the following section. We analyze each of the best-found trading rules individually, including tick size, priority rules and multi-price rules, transparency and the intervention of market maker.

Simulation	Volume	Fraction specialist (θ)	Fraction revised (α)	Price adjustment (λ)	Fraction informed (γ)	Tick size (t)	Multi-price rule	Priority rule
1	17.509	0.925	0.900	0.186	0.024	5.140	Lowest	Time
2	17.811	0.988	0.951	0.150	0.034	5.300	Lowest	Time
3	18.072	0.991	0.671	0.234	0.015	4.151	Lowest	Reverse time
4	18.147	0.977	0.691	0.619	0.015	4.809	Lowest	Size
5	18.919	0.987	0.432	0.456	0.035	4.064	Lowest	Reverse time
6	19.001	0.959	0.823	0.129	0.024	3.603	Highest	Size
7	19.695	0.948	0.437	0.814	0.012	3.359	Highest	Random
8	19.837	0.999	0.995	0.909	0.021	3.696	Highest	Size
9	20.001	0.955	0.743	0.024	0.011	3.528	Highest	Size
10	20.027	0.983	0.730	0.814	0.007	3.919	Lowest	Time
11	20.549	0.991	0.659	0.924	0.008	2.767	Highest	Time
12	20.557	0.998	0.764	0.883	0.012	3.809	Lowest	Time
13	20.875	0.799	0.693	0.450	0.018	2.951	Highest	Reverse time
14	21.424	0.841	0.949	0.435	0.036	0.965	Closest	Size
15	22.023	0.996	0.889	0.323	0.031	0.838	Closest	Random
16	23.396	0.978	0.965	0.886	0.015	0.994	Closest	Size
17	24.178	0.989	0.774	0.602	0.020	0.196	Closest	Time
18	24.278	0.991	0.996	0.921	0.032	0.128	Closest	Time
19	25.463	0.991	0.712	0.434	0.028	0.485	Closest	Reverse time
20	25.974	0.915	0.683	0.466	0.032	0.057	Closest	Time
Mean	20.887	0.935	0.773	0.533	0.022	2.738	Closest/*	Time*
StDev	2.565	0.147	0.164	0.301	0.010	1.792	Lowest**	

Tab. 5.2: Values of the maximized trading volume and the best-found trading rules for each optimization simulation * The mode is shown instead of the mean for priority rules and multi-price rules. ** Two modes exist.

5.2.1.1 Tick size

As revealed in table 5.2, the values of tick size fluctuate widely between the minimum value of 0.057 and the maximum value of 5.300. The mean value of the best-found tick size from the 20 simulations is 2.738, which is significantly higher than the tick size of one-sixteenth dollar previously applied on the NYSE. Therefore, not only the small tick sizes, which are normally applied in most of the stock exchanges, can generate high trading volume; large tick sizes can also be employed to maximize the trading volume.
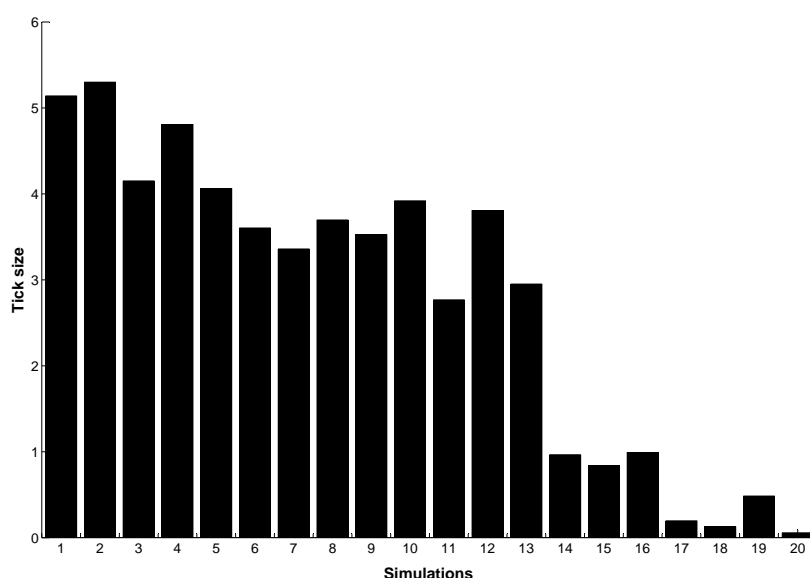

Fig. 5.1: Simulation results of the best-found tick size for maximising trading volume

In figure 5.1, we report the simulation results of the best-found tick size for maximizing trading volume. Since the results are sorted based on the trading volume, the first observed tick size in the figure corresponds to the lowest volume generated, while the last tick size corresponds to the highest trading volume. It is shown in figure 5.1 that the best-found tick sizes disperse in two regions. When trading volumes are comparatively higher (trading volumes of simulations 14 to 20 in table 5.2), the corresponding tick sizes are smaller than one. While volume are relatively lower (trading volumes of simulations

1 to 13 in table 5.2), the tick sizes are, on the contrary, very high. The result is suggesting that small tick size leads to higher trading volume, and large tick size generates lower trading volume.

While tick size has been set small, any two submitted quotes would be very close to each other. This leads to the case that the submitted orders will overlap more easily, which would then help to match the buy and sell orders. In other words, the closer the submitted quotes, the easier buy and sell orders would be matched up. It naturally increases the chances of trading when new incoming orders are submitted. Consequently, it leads to an increase in trading volume. On the contrary, if the tick size has been set too large, the gap between order quotes from different traders would be very wide, which may increase the difficulties of matching buy and sell orders, and hence decrease the chances of trading. As a result of this, the trading volume would decrease. Therefore, our results suggest that small tick sizes lead to higher trading volume, and large tick sizes tend to generate lower trading volume.

This result is also consistent with the literature, e.g., see Harris (1994), Allaudeen and Eric (1998), and Chan and yang Hwang (1998). Harris (1994) find that the trading volume would be enhanced with a decreased tick size. Allaudeen and Eric (1998) suggests that trading volume is more likely to be increased with a reduced tick size if the stock is actively traded. As discussed by Allaudeen and Eric (1998), this result is consistent with the price resolution hypothesis. In addition, they explain that "Direct trading costs are lower when the tick size is reduced and may, as a result, increase trading volume. However, higher negotiation costs may dampen the increase in trading volume due to lower direct trading costs. Hence, a reduction in tick size may not always lead to increased volume."

Observation 1: From table 5.2 and figure 5.1, it follows that smaller tick size will lead to a higher trading volume and vice versa.

5.2.1.2 Priority rules and multi-price rules

Based on the experimental results shown in table 5.2, we do not observe a specific priority rule that maximize the trading volume. Moreover, there is no clear pattern relating the priority rules to the trading volume or other rules selection. All the potential priority rules, except for the rule of pro-rata, can be applied to maximize the trading volume. However, the time priority rule is the most commonly obtained rule from the simulations. This is also the most common rule used in real markets.
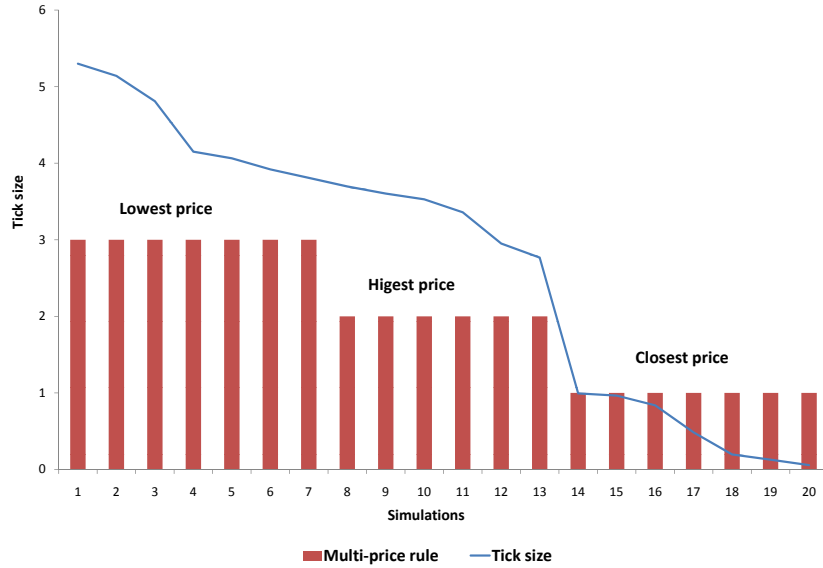

Fig. 5.2: the best-found tick size and multi-price rule for maximizing trading volume

The bar charts represent the different multi-price rules. They do not refer to any values. Different heights are assigned to different bars to indicate the difference of rules.

For multi-price rules, three potential rules could be used to achieve the desirable result, which are to choose the highest or lowest price, or the one closest to the previous price. Among these, the most frequently obtained multi-price rules are the lowest and closest price rules. In addition, we find a connection between the tick size and the multiple price rule, see figure 5.2. For instance, when the closest multi-price rule is obtained to maximize the trading volume, the best-found tick size is smaller than 1; with the multi-price rule that

highest price is selected, the tick size increases to the range of 2.77 to 3.70 (see table 5.2 for the data); and it further increases with the multi-price rule that the lowest price is selected. This result indicates a trade-off between the size of the bid-ask spread and the multi-price rule that would become relevant should the bid-ask spread be used as an alternative or additional performance indicator for the market.

From our results, when small tick size is found to be desirable to maximize the trading volume, the best-found multi-price rule is the closest price rule. When tick size is small, the submitted quotes are easily to get closer together within a comparatively small range. In this case, in order to enhance the trading volume, sell order prices should be around a price at which more trades would occur[1]. With the closest price rule, we select the closest potential price to the previous market price as the market price[2]. Sell order prices in the next round of trading, which depends on the market price in this round, would then be close to not only the market price at which the trading volume is maximal in the previous round, but also the market price in the two rounds before. Thus, the probability of matching with buy orders would be much higher, and the trading volume would be enhanced as well. Orders in the middle part of the order book would be executed, leaving orders with extreme prices remaining in the order book for future trade.

On the other hand, when large tick sizes are desirable to maximize the trading volume, the best-found multi-price rules are either the highest price rule or the lowest price rule. While tick size is large, the submitted quotes are further

---

[1] Buy order prices do not move much from time to time as they are randomly chosen from the distribution based on the fundamental price, whereas sell order prices always change according to the last purchasing price, which is the market price in this auction market, see section 4.2. Thus, the sell order prices are an important factor in influencing the trading volume because if they can easily match with bid prices the trading volume will be higher, otherwise it is lower. Sell order prices are hence our main concern here.

[2] The market price is always chosen at which the trading volume is maximal, see section 4.2.4. In some cases, there might be multiple prices at which the trading volume shows the same maximal value. Thus, multi-price rules are adopted to determine which of the potential prices to be the market price. In this case, it is the price which is the closest to the previous market price.

apart from each other, and the range of the prices are comparatively larger to the case when tick size is small. In this case, if we still use the closest price rule and focus on the central part of the order book, it would make less matching between buy and sell orders, which in turn reduces trading volume. As large tick sizes pushes submitted quotes away from each other, there would be more orders waiting to be traded at both ends of the order book rather than in the center of the order book. Thus, more trading would be executed if we focus on both ends of the order book. With the highest/lowest price rule, the sell orders (with prices based on the previous market price: either the highest potential price or the lowest potential price) in the next round are moved away from the middle to some extreme areas in the order book where orders have been left unfilled from previous trading rounds. Hence, apart from the matching between new orders, the remaining orders from previous rounds can also be matched and traded with the new submitted orders. Therefore, the trading volume is enhanced.

From figure 5.2, we also observe that when the best-found tick sizes are the largest, the lowest multi-price rule is desirable to maximize the trading volume; while the best-found tick sizes are comparatively lower, the highest price rule is found to be the best multi-price rule. If the tick size is very large and the highest price rule is used, the sell order prices might have too much movement, which then may lead very high prices and make it difficult to find buy orders to match, i.e., few buyers are willing to accept and trade at such high prices. This asymmetry of the orders would consequently reduce the trading volume. However, if the lowest price rule is adopted, though the large tick sizes lead to a large spread of prices, the sell order prices, in this case, are comparatively lower, and hence there would be comparatively more matching between buy and sell orders than the case that the highest price rule is adopted. Once the tick size reduces, the lowest price rule could only make a small movement of the ask prices that not much orders could be matched. Instead, the highest price rule should be employed, whereas the movement is large enough to induce more orders to be matched and hence the trading volume is maximized.

Observation 2: As apparent from figure 5.2, a relation between tick sizes and multi-price rules emerges while obtaining the best market structure.
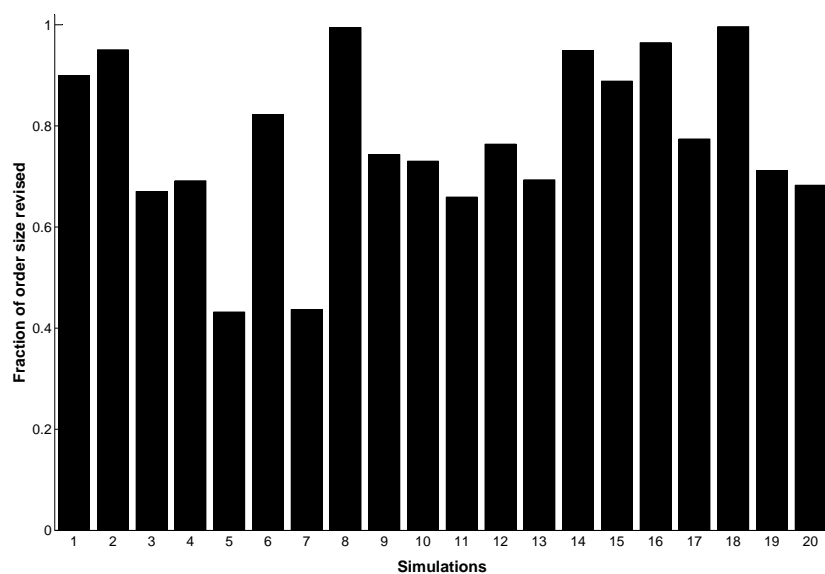
5.2.1.3 Transparency

As discussed in section 4.2.4, in a transparent market, traders can observe information on the order book and revise their own order size based on the market information of order imbalance. With respect to market transparency, the best-found market structure indicates that in order to maximize trading volume a very small fraction of 2.2%, on average, of traders should be able to access information on the order book and react to any ensuing order imbalance if the trades are to happen immediately. The parameter values, with a small dispersion, as shown in figure 5.3(a), does not change much from the mean in different simulations.

Our results show that the low transparency, as represented by a small fraction of traders are able to access information on the order imbalance, does not necessarily reduce trading volume but rather the reverse (i.e. less traders informed on the order book information, more orders would be traded). Since transparency allows traders to observe information about the imbalances between current market supply and demand from the order book, as well as to revise their orders before the trading is taking place, the higher the level of transparency the wider the discrepancy between the revised cumulative buy and sell order sizes (i.e. if more traders revise their orders accordingly based on their observation on the order imbalance, the cumulative order sizes for buy and sell would be further apart from each other). For instance, if the order imbalance is greater than zero - that is when buy order sizes is greater than sell order sizes at the transaction price, all informed traders would revise down(up) their buy(sell) order sizes not only at the transaction prices but also at all the other price levels, and hence increases the gap between the cumulative order sizes for buy and sell. Although the order revision may reduce the order imbalance between the supply and demand from the order book, the level of trading volume is only dependent on the number of the
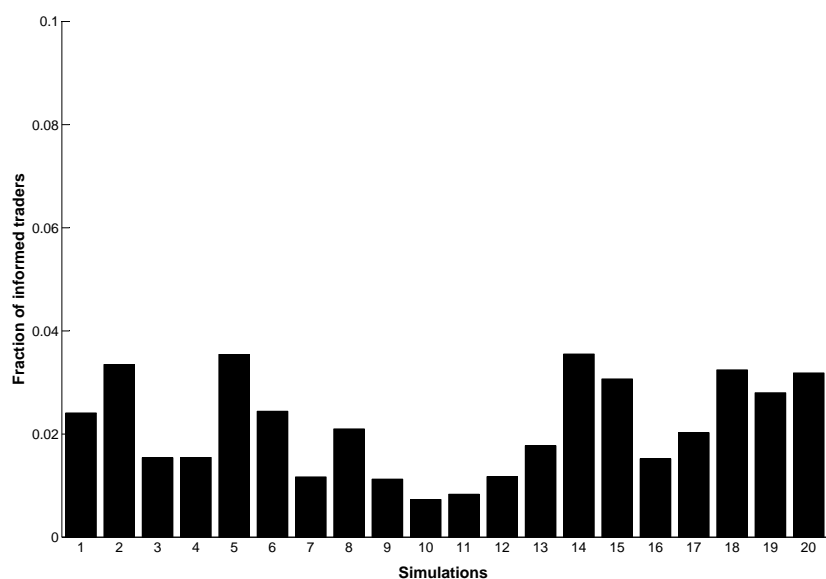
(a) the best-found fraction of informed traders

(b) the best-found fraction of order size revised

Fig. 5.3: Simulation results of market transparency for maximizing trading volume

most possible matched orders between the supply and demand sides of the order book, which are obtained by using the cumulative buy and sell order sizes. As a result of the large gap between the revised order sizes induced by

high level of transparency, the trading volume would reduce to a level lower than in the case where fewer traders are informed on the order imbalance.

This result is consistent with the literature. As demonstrated by Goldstein et al. (2007), with increased transparency, "Measures of trading activity, such as daily trading volume and number of transactions per day, show no relative increase, indicating that increased transparency does not lead to greater trading interest". In addition, Aitken et al. (2001) confirm that "the enhancement in pre-trade transparency, through tightening the undisclosed order regulation in October 1994, resulted in a significant decline in trading volume."

Moreover, based on the results, the informed traders are suggested to revise their order sizes with the adjustment of, on average, 77.3% of the order imbalance. Except for two values (of simulation 5 and 7), all other resulting parameter values are above 60%, see figure 5.3(b). Intuitively, the impact of the revisions would be bigger if a large fraction of order size has been revised. Yet, our previous findings suggest that very large would deteriorate trading volume. However, when taking into account our previous results that on average only 2.2% of the investors are informed (i.e. only a small fraction of investors would access the information and revise their order sizes accordingly), the overall impact of the order size revision is significantly diminished. In other words, although the best-found order size fraction to be revised is large (77.3%), the impact of this revision is small for the reason that only limited number of investors who are able to revise their orders.

Observation 3: Following figure 5.3(a), it is found that to increase trading activities markets should provide less transparency.

5.2.1.4   The intervention of market maker

Lastly, we analyze how trading volume can be maximized with the intervention of market maker. As shown in table 5.2 and figure 5.4(a) the average value of the fraction of order imbalances traded by market maker from the
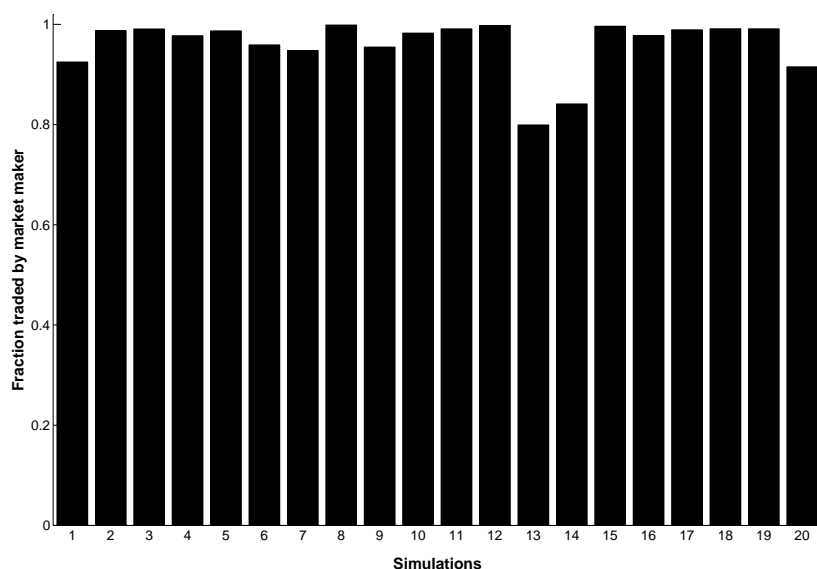
(a) the best-found fraction of order imbalance traded by market maker for maximizing trading volume
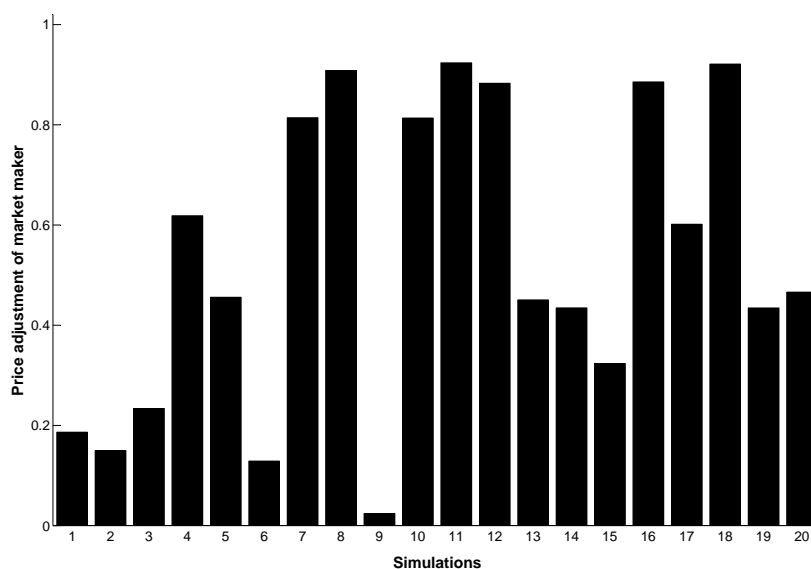
(b) the best-found price adjustment of market maker

Fig. 5.4: Simulation results of intervention of market maker for maximizing trading volume

experimental simulations is very high (at 93.5%), while most of the values are around 90%. It thus follows that market maker should be more active to

trade in the market in order to increase the trading volume. This is compatible with the common belief and common knowledge. If market maker trades more in the market, more unfilled orders can be traded and cleared, which in turn increase trading volume.

Similar to the priority rules, we do not observe a clear pattern for the desirable value of price adjustment of market maker to reach the maximal trading volume, i.e. we observe maximum trading volume under either high or low price settings by the market maker. This is due to the fact that it is very difficult to identify a particular value of the price adjustment of market maker. The transaction price submitted by the specialist is formed using equation 4.8, where λ is averaged at 0.533, with a quite broad range from 0.024 to 0.924, see table 5.2 and figure 5.4(b). Consequently, as our results suggested the price setting of the market maker does not affect the trading volume. This may because that when the market maker is intervening the market by trading a part of the order imbalance (either as a buyer or a seller), his price setting depends on its own inventory levels rather than the current order imbalance, therefore the impact of the price setting by the market makers on the trading volume is vague.

Observation 4: From table 5.2 and figure 5.2.1.4, it follow that the intervention of market maker stimulates trading activities and leads to a higher trading volume.

5.2.1.5 Summary of results

The main results of the optimization of market structure with the fitness function of trading volume can be summarized as follows. In order to achieve the maximal volume, the stock exchange should implement a tick size of either from the range of 0 to 1 or from the range of 2.77 to 5.30, depending the selection of multi-price rule; allow only a very small proportion (about 2%) of traders access to information in the order book, but a high fraction of order imbalance should be traded by market maker. The use of priority rules,

on the other hand, does not show any correlation to the trading volume. Our results are consistent with real market observations, except for the selection of tick size. We can therefore conclude that the main effect on trading volume arises from the tick size, multi-price rule, access to the order book, and intervention of market maker.

5.3 Summary of single-objective optimization experiments results

We have analyzed the best-found trading rules for each optimization experiment with a different objective function (see section 5.2.1 and A. This section summarizes the main results found from the experiments.

Apart from the small ticks used in real markets, significantly large tick sizes have also been found to be desirable in most of the optimization cases investigated in this research. Our results suggest that smaller tick sizes lead to a higher trading volume while larger tick sizes reduce volume, see Observation 1 on page 138. In addition, a higher tick size, instead of a smaller tick, should be implemented in order to minimize the bid-ask spread, which differs from previous research, see Appendix A.

A special connection between the best-found tick sizes and multi-price rules has been observed in all the experiments with various fitness functions, see Observation 2. For example, when the best-found tick size of maximizing trading volume is smaller than 1, the closest multi-price rule is selected; when the tick size increases to the range of 2.77 to 3.70, the multi-price rule that highest price is obtained; and if the tick size further increases, the multi-price rule that selecting the lowest price is employed.

We do not observe a unique combination of priority rule and multi-price rule. The time priority rule and the closest multi-price rule are the most frequently obtained rules. However, all the other rules, except for the rule of pro-rata,

have been attained in the optimization experiments. Similarly, as the result for the market transparency and the extend of market maker intervention vary in wide ranges, it is hard to identify a particular desirable value.

Finally, we find that the results of the best-found trading rules differ significantly with different objective functions, e.g., the best-found combination of trading rules that maximizes the trading volume is unable to minimize the bid-ask spread. For market transparency, it has been found that in a less transparent market trading volume is much higher, see Observation 3. However, it is not always true that a low degree of market transparency would achieve the best market structure with different objective functions, e.g., the influence of market transparency are vague on trading price and volatility, see Appendix A. Similarly, while intervention of market maker stimulates trading activities and increases trading volume, see Observation 4, it, however, suggests to be maintained at a lower level to minimize market volatility, see Appendix A. Given that market participants often value different aspects of an exchange and use different performance measures simultaneously, we need to employ multiple objectives when optimizing the market structure in order to consider the various concerns of market participants.

6. Optimization of market structures with multiple objective functions

The single-objective optimization of market structures has been conducted using trading-volume, trading-price, bid-ask spread or return volatility as the performance measurement. In this section, we extend the framework to a multi-objective setting to evaluate how any conflicts between different interests in market characteristics might be resolved. For simplicity this research will only consider cases consisting of two objective functions, e.g., maximizing trading volume while minimizing the bid-ask spread. Though the developed framework is also valid for higher dimension cases, the use of only two objectives ensures a straightforward visualization of the results obtained.

6.1 Population-based incremental learning for multiobjective optimization

Applications of multi-objective optimization using the PBIL algorithm are rare thus far, e.g., Bureerat and Sriworamas (2007), and this paper is one of the few applications in this field. Instead of a single probability vector used in the PBIL for single objective optimization, a probability matrix containing numerous probability vectors is created in the MOPBIL. Each probability vector is used to generate a sub-population, and updated, similar to the

single-objective optimization, using equation 5.1. In addition, the probabilities are subject to mutation at a mutation rate ξ and the actually chosen probability π_t^* will be

$$(6.1) \qquad\qquad \pi_t^* = (1 - \xi)\pi_t + \xi\varepsilon$$

with $\varepsilon \sim U[-1; 1]$ and π_t^* restricted between 1 and 0. This algorithm is capable to maintain diversity in search as the same probability vector could generate distinct populations.

To employ the multi-objective PBIL algorithm, we first create a probability vector specifying the probability of each bit position having the value one. All the initial value for the PBIL process is determined randomly, i.e., the initial probability of each bit position is considered as 0.5. With this probability vector a population is produced according with its values. In each generation we determine the fitness of each potential solution and the best set in the current generation is selected and used to update the probability vector for the subsequent generation using (5.1) and (6.1).

In each time step we determine 100 different parameter constellations using π_t^* and then determine the best performing parameter constellation from these 100 different market simulations that then makes $\widehat{v}$. Each trading rule is coded by the same method we applied in the single-objective optimization.

As is common with multi-objective optimization, we do not observe an easy convergence of results (even after 5,000 generations no convergence towards a clearly identifiable Pareto-efficient frontier was observed), see section 3.3.5 for examples and discussion of these issues. For this reason we run the optimization for 500 generations and use the entire population of the resulting final generation to analyze our results. Investigating the evolution of the Pareto-efficient frontier we observe that using more generations does not improve the solutions presented here as the Pareto-efficient frontier does not move any further. During the last few hundred generations we only observe an oscillation of the efficient frontier without any significant improvements. The number of generations used is about 4 times the length it took the

single-objective optimization in section 5.1 to converge and should therefore represent an adequate time length for the evolutionary algorithm to evolve. As reviewed in section 3.3.5, it is common in multi-objective optimizations to determine a fixed number of generations after which the optimizations is stopped. Thus, our approach is very much in line with the literature.

To obtain the optimal market structure with multiple performance functions, we apply the same parameter values used in the single objective optimization simulations, see section 5.1. However, instead of 200 generations for the single-objective optimization, the multi-objective PBIL optimization is conducted over 500 generations with learning rate of 0.2 and a mutation rate of 0.01 for the reason discussed in section 3.3.5. We repeat the multi-objective optimization 30 times, which are 10 times more for the single-objective optimization simulations, with the same parameter constellation to obtain more data and reduce the amount of noise remaining from the lack of convergence.

6.2   Evaluation of computer experiments

Four experiments are conducted to investigate the optimal market structures using multi-objective PBIL algorithm. Each experiment seeks to obtain the best combination of trading rules which can optimize two specific objective functions simultaneously, e.g., maximizing trading volume while minimizing bid-ask spread; maximizing trading price and minimizing bid-ask spread; maximizing trading volume and minimizing return volatility, and finally minimizing bid-ask spread while minimizing return volatility. The experimental results of the best-found market structure with the objective functions of maximizing trading volume and minimizing bid-ask spread are presented in the following part. The other experimental results are discussed in Appendix B.

6.2.1   Trading volume and bid-ask spread

First of all, we only consider two objective functions namely: trading volume and bid-ask spread. The aim is to find the Pareto front which consists of the best combination of trading rules that maximize trading volume and minimize bid-ask spread. In figure 6.1 we show the maximized trading volume and minimized bid-ask spread of the final generation for the entire population for all 30 runs of our computer experiments restricted to the area close to the Pareto-efficient frontier. All the points in the figure are solutions found from the simulations. We have identified nine markets that approximately determine the Pareto-efficient frontier, these points are identified as large points and associated with numbers. We observe a trade-off between the maximal trading volume and the minimal spread, the approximate location of the Pareto-efficient frontier is sketched by the line to the lower right. This figure clearly shows that a small spread will be associated with a low trading volume while a high trading volume will necessitate a large spread. As shown in figure 6.1, market 1 has the smallest trading volume and bid-ask spread, while market 9 has the highest trading volume and spread. This positive relationship is very pronounced; such results are commonly found for Pareto-efficient frontiers not only for multi-objective optimizations but also portfolio selection theory. The desirable combination between trading volume and spread will depend on the preferences of the decision-maker and how he values the importance of these two aspects; if the preferences are such that concerns about the bid-ask spread dominates, small changes to these preferences will result in a very different location of the optimal market structure on the Pareto-efficient frontier.

Observation 5: As shown in figure 6.1, a trade-off between the maximal trading volume and minimal bid-ask spread has been observed. A small trading volume is always associated with a small spread, while a higher trading volume is associated with a higher spread.
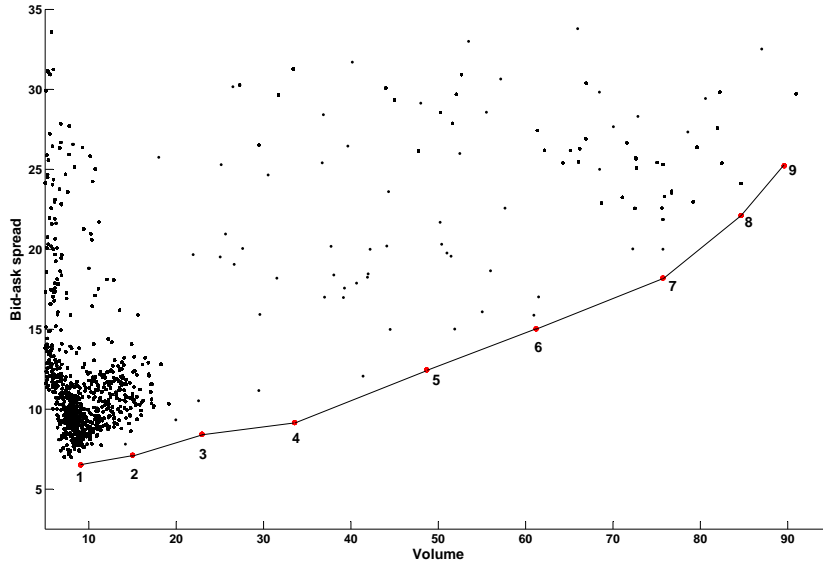
Fig. 6.1: Market performance of maximizing trading volume and minimizing bid-ask spread after 500 generations

6.2.1.1 Tick size

Figure 6.2 presents the best-found tick size of the obtained nine market structures that approximate the pareto-efficient frontier. It shows that when concerns about the bid-ask spread are dominating (e.g., in market 1 where spread is the minimal), a large tick size should be chosen and a small tick size if concerns about trading volume are more important (e.g., in market 9 where trading volume is the highest). The results are consistent with the finding in the the single-objective optimization experiments for maximizing the trading volume and minimizing bid-ask spread: when maximizing trading volume a small tick size would be desirable; when minimizing inside spread a large tick size would be more favorable.

This finding is analogous to that discovered in the single-objective optimization. As explained in appendix A.1, when tick size is large, quotes of many orders might be centered around a few ticks. As a result, the bid-ask spread after trading would be narrowed. On the other hand, a smaller tick size leads to order quotes being closer together. Hence, an incoming order is more likely

to be matched and executed rather than remaining stored in the order book, awaiting another order to be submitted before trade occurs. Consequently, the trading volume increases with the lower tick size, see section 5.2.1.1.

Observation 6: The best-found tick size found in the multi-objective optimization experiment of maximizing trading volume and minimizing bid-ask spread follows the similar trend that has been observed in the single-objective optimization experiments. As shown in figure 6.2, a large tick size is preferred when concerns about spread are dominating, and a small tick size should be used if concerns about trading volume are dominating.
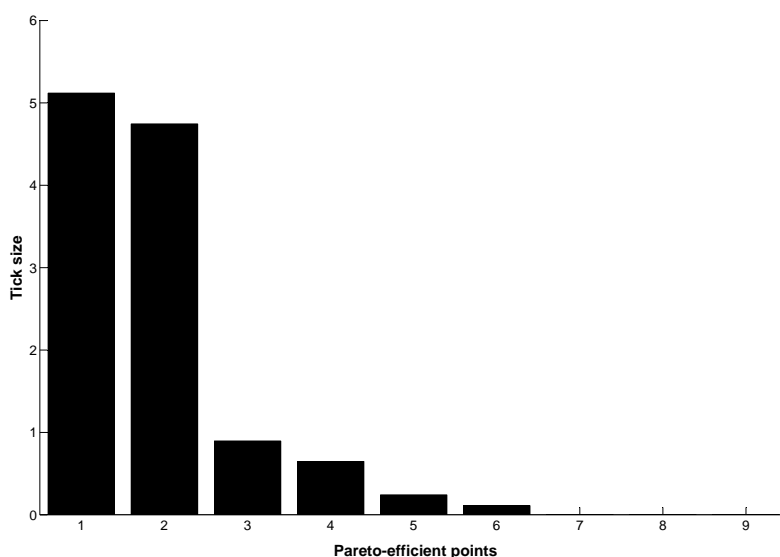
Fig. 6.2: the best-found tick sizes for maximizing trading volume and minimizing spread

6.2.1.2 Multi-price rules and priority rules

The best-found multi-price rule is to choose the nearest price to the previous price, except for the first two markets where the best-found multi-price rule is to select the lowest price in case of multiple prices, see figure 6.3. Again, the relation between the selection of multi-price rules and the best-found tick
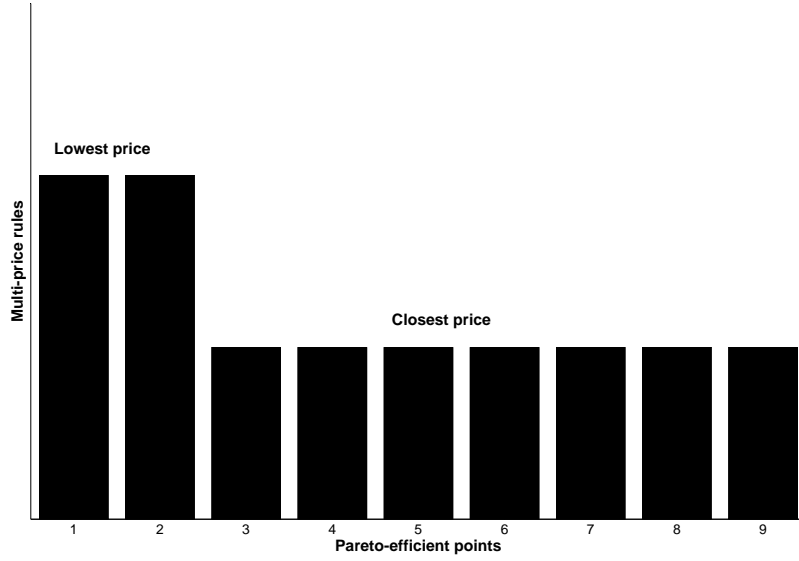
Fig. 6.3: the best-found multi-price rules for maximizing trading volume and minimizing spread

sizes, which has been observed in the single-objective optimization simulations in section 5.2.1.2, has also been found in this experiment, see figure 6.2 and 6.3. When the best-found tick size is less than 1, the closest multi-price rule is found to be the desirable rule to maximize trading volume and minimize the inside spread. When the tick size is larger (around 5), the lowest multi-price rule is the best-found rule. Order prices tend to close together when tick size is small. In order to increase the volume, sell order prices should be around the price at which maximal trades occur. The closest price rule is thus selected as it choose the closest potential price to the previous trading price as the market price in this round. With this rule, sell order prices in the next round would be close to not only this round's market price but also the previous round's market price at which the maximal trading volume exists. Consequently, the probability of matching between buy and sell orders are higher. On the other hand, when tick size is large, the submitted quotes spread widely in the order book, which causes more orders waiting to be traded at both ends of the order book. Hence, the lowest price rule, which allows the sell orders moving away from the middle to some extreme areas

in the order book, is selected (for detailed explanations, see section 5.2.1.2).

Three priority rules are found capable to maximize the trading volume and minimize the inside spread, see figure 6.4. The best-found rules include selecting the orders randomly, or based on their submission time (i.e., first in first out), or based on the reverse time selection rule (i.e., last in first out), of which the time priority is always selected especially when trading volume is considered as the main aim by the decision maker.

Observation 7: The relation between the best-found tick size and the selection of multi-price rules has also been observed in the multi-objective optimization experiments.
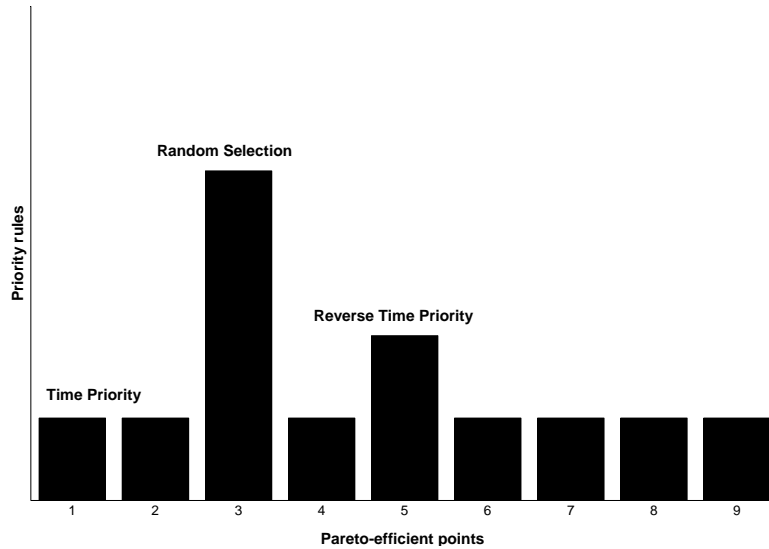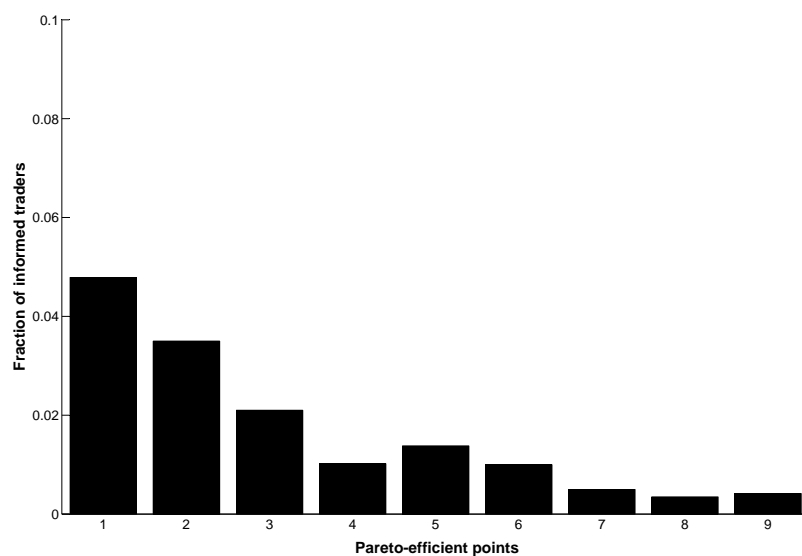

Fig. 6.4: the best-found priority rules for maximizing trading volume and minimizing spread

6.2.1.3 Transparency

As displayed in figure 6.5(a), the results indicate that a very small fraction of approximately 2% of traders are suggested to access information on the order book and revise their order size based on the order imbalance. In

(a) the best-found fraction of informed traders for maximizing trading volume and minimizing spread
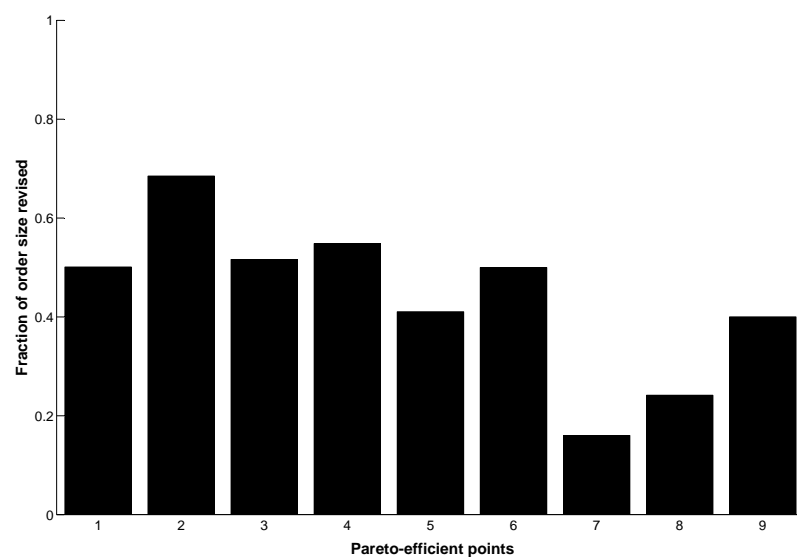
(b) the best-found fraction of order size revised for maximizing trading volume and minimizing spread

Fig. 6.5: Simulation results of market transparency for maximizing trading volume and minimizing spread

general, we observe that a decreasing fraction of informed traders leads to a higher trading volume (and associated with a larger bid-ask spread). For instance, in market 1 where trading volume is very low compare to the other

8 markets, the best-found fraction of informed trader is very high at 4.79%; when trading volume increases to 89.6 in market 9, the best-found fraction decreases to 1.56%.

This is consistent with the single-objective simulation results of maximizing trading volume that trading volume would increase if the fraction of informed traders maintained at a low level. This is because that if more traders are informed, the discrepancy between the revised cumulative buy and sell order sizes would be wider as all informed traders would revise their order sizes according to the order imbalance. This would then lead to a smaller trading volume as the matching orders between the supply and demand are reducing, see 5.2.1.3 for more detailed explanations.

Similar to the results of single-objective optimizations of maximizing trading volume and minimizing bid-ask spread, no clear pattern has been observed for the fraction of the order imbalance revised by informed traders. Figure 6.5(b) reveals that the best-found fraction revised by informed traders seems to be independent of the location on the Pareto-efficient frontier at approximately 44%, and it fluctuates between 10% to 0.7%.

Observation 8: Following figure 6.5(a), it is found that a decreasing fraction of informed traders stimulates trading volume and enlarges bid-ask spread.

6.2.1.4   The intervention of market maker

We notice that unless the decision maker puts great emphasis on a large trading volume at the expense of the spread, the fraction of order imbalance traded by market maker should be very small as can be observed from figure 6.6(a). It increases dramatically with more concerns on trading volume. The intervention of a market maker obviously increases the trading volume as he trades otherwise unfilled orders. However, results are ambiguous for the price adjustment of market maker, here no clear pattern can be identified for different points along the Pareto-efficient frontier, but, all the values are very low, less than 0.06.
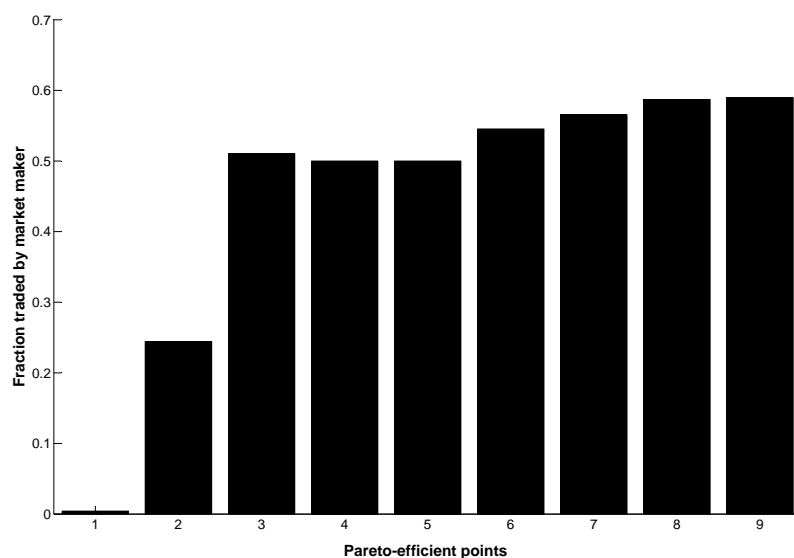
(a) the best-found fraction of order imbalance traded by market maker
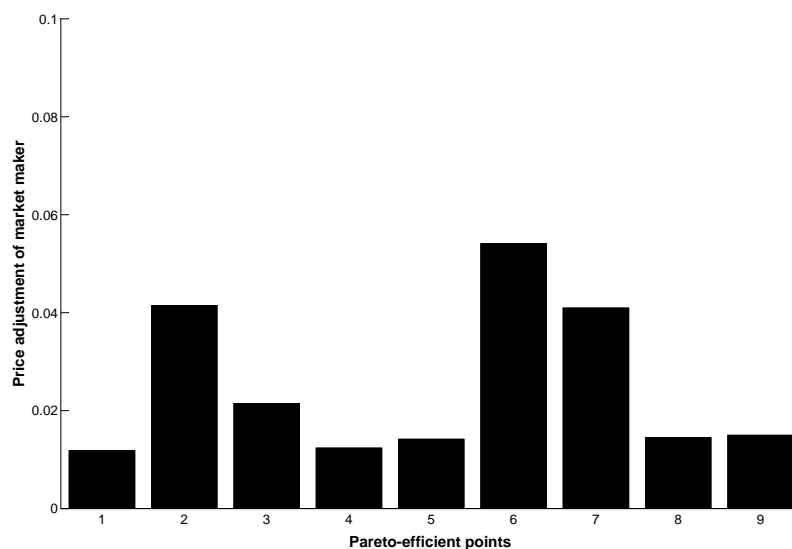
(b) the best-found price adjustment of market maker

Fig. 6.6: Simulation results of intervention of market maker for maximizing trading volume and minimizing spread

These results are much in agreement with the experimental results of maximizing trading volume. The intervention of market maker do increase the trading volume as he trades orders which cannot be filled in the market. On the other hand, the price setting of market maker, which correlates to

the market maker's inventory levels instead of the current order imbalance, shows no clear impact on the trading volume.

Observation 9: As shown in figure 6.6(a), the best-found fraction of imbalance traded by market maker is comparatively higher when trading volume is the main concern, instead of the bid-ask spread.

6.2.1.5   Summary of results

We have analyzed experimental results of maximizing trading volume and minimizing bid-ask spread. A trade-off has been found between the two objective functions: while trading volume is higher, the corresponding bid-ask spread is lower, and vice versa. The main results from our computer experiments as outlined above are that for solutions requiring a larger trading volume (and associated higher bid-ask spreads) smaller tick sizes are required, the best-found degree of market maker intervention should be high, while the best-found fraction of informed traders should be maintained at low level. Time priority rule and the multi-price rule of choosing the nearest price from the previous price are the most applied rules when concerns about the trading volume dominates. The relation between the tick size and multi-price rule has again been observed in the multi-objective optimization simulations. We also notice that most of the results obtained in this section are much consistent with the results obtained in the single-objective optimization simulations of maximizing trading volume and minimizing bid-ask spread.

6.3   Summary of multi-objective optimization experiments results

The main results from our computer experiments of multi-objective optimization as outlined above are that only large tick sizes are required for

maximizing price and minimizing bid-ask spread; for other objective functions, the best-found tick sizes are distributed widely from 0 to 6. Our result that a larger tick size is required is not in accordance with the tick sizes used in real markets, but it is consistent with findings in single objective optimizations in section 5.1. Moreover, most of the priority rules and multi-price rules have been attained to achieve the best market structure, except for the farthest multi-price rule and the rule of selecting the price based on order imbalances; and pro-rate and size priority rules. And, the most often obtained rules are the lowest multi-price rule and time priority. For the market transparency, we find that, in general, only small fraction of traders should be informed in order to get the best market performance with different fitness functions. This result is the only one that remains unchanged in various experiments. Besides, the best-found intervention of market maker shows ambiguous results. Finally, analogous to the previous results, we do not find a particular combination of trading rules that could be applied to optimize the market structure with different pairs of objective functions.

Comparing these results with those in section 5.1 in which the same model and optimization technique are used for single objective function, we find that our results here are largely consistent with those obtained in the last section. Most of the best-found trading rules obtained here are much similar to those obtained in the single-objective optimization, or observe the same trend to achieve the desirable market structure, where many were found to be very variable when running different experiments. For example, similar to the results from single-objective optimization, most of the best-found tick sizes attained here remain significantly higher than the tick size employed in real markets. In addition, the connection between the best-found tick sizes and multi-price rule selection observed in single-objective optimization experiments has also been found in the multi-objective optimization simulations, see Observation 7. Also, the most frequently obtained priority rules is the time priority in both single and multiple objective function optimizations. This rule has been applied in many real markets. Different from the single-objective optimization, the market transparency shows a clear trend in

the multi-objective optimization in which a small number of informed traders
are much preferred.

7. Conclusions and discussion

We investigated evolutionary mechanism design of stock markets though the use of automatons and evolutionary optimization algorithm. This final chapter concludes the work by summarizing the research, considering the main contributions, and discussing its limitations and suggestions for future work.

7.1 Research summary

Stock exchanges today face increasing competition due to the powerful trends of migration of trading and listings abroad. To attract more companies for listing their stocks and investors for trading, exchanges require to offer an efficient, fair and competitive trading environment. They have begun to improve their market structures or change their trading characteristics to become more competitive. Market design is therefore taking on new importance that markets need to compete (O'Hara (2001)).

To provide insights into the optimal market structure, this research explores the evolutionary design of call markets in which orders of traders batched, i.e., order submissions collected over a short period of time and then executed in a single trade. The most common call markets are financial markets which provide the background to this research project. This work relates the market microstructure theory to the mechanism design problems by investigating a wide range of potential trading rules, including tick size, priority rules, multi-price rules, market transparency, and intervention of market maker, commonly found in financial markets to determine the optimal combination of these rules to enhance the market performance.

In order to overcome the drawbacks of conventional models in the market microstructure literature, we develop an agent-based model featuring near-zero-intelligence traders. In this approach, traders follow very simple trading rules which does not assume rational behavior or any other optimizing rule. Introducing heterogeneous traders along various aspects also enable us to generate realistic market characteristics. The main objective of this research is to investigate optimal market structure. It has been proved that the primary cause of the high allocative efficiency of auction markets and the convergence of transaction prices to equilibrium levels is the market discipline imposed on traders; intelligence or profit motivation is not necessary (see Gode and Sunder (1993)). The use of zero-intelligent traders in this research allows us to focus on the influence the market structure has on the outcomes without any impacts from traders' behaviors. In addition, it has been found that the use of zero-intelligence traders closely mimics stock market and generates realistic and efficient results, see section 2.2.3. In contrast to our approach, the results of conventional models are often dominated by the behavior of the agents and it proves difficult to distinguish their impact from that of the market structure.

Having developed an appropriate model, we investigate the credibility of our model by reproducing stylized facts through the trading behavior of near-zero intelligence traders. Based on computer simulations, we confirm that our model is capable to reproduce some of the most important stylized facts found in financial markets, including the long memory of return volatility; fat tails of the return distribution; and multi-scaling of returns. We also study the stylized facts under various market structures. It has been found that the properties observed do not vary significantly with the trading rules applied; only with a small number of trading rules did we observe a change in properties. This outcome is in agreement with the observation that the stylized facts are observed in most asset markets, regardless of trading rules applied.

After confirming the credibility of our model, exploring the optimal market structures, which is the main objective of this research, is then conducted evo-

lutionary using population-based incremental learning as an efficient search algorithm in this complex problem. Market participants often value different aspects of an exchange and will thus likely to use different performance measures, requiring the use of different objective functions to find the optimal market structure. Therefore, a wide range of objective functions, which are commonly believed as of traders' interests, is used in this research, including maximizing trading price or trading volume, minimizing bid-ask spread or volatility of returns.

We aim to optimize the market structure based on both single objective function and multiple objective functions. From the experiments with single objective, we find a quick convergence of the trading rules towards their best-found combination. We have analyzed the combination of trading rules for each experiment with various fitness function. And, the best-found set of trading rules differ significantly if objective function changes.

This framework is then extended to a multi-objective setting to evaluate how any conflicts between different interests in market characteristics might be resolved. In order to identify the best possible market structure we use multi-objective population-based incremental learning algorithm to derive the Pareto-optimal set of market structures. In order to choose the market structure that should be implemented, the market structures of those markets close to the Pareto-efficient frontier have been investigated further. Based on the results, we can determine the desirable market structure for each objective function, which would have direct consequences for the optimal design of financial markets in terms of each objective function. Our results could inform any market reforms considered by stock, bond or derivatives markets.

The results of single-objective optimization show that significantly large tick sizes contribute to most of the optimization cases, except for the case of maximizing trading volume. In addition, the tick size always correlated with the selection of multi-price rules. Though there is no particular combination of priority rule and multi-price rule are desirable, the time priority rule and the

closest multi-price rule are the most frequently obtained rules. The market transparency and the extend of market maker intervention show ambiguous results as their representative parameter values change in a wide range. Finally, we do not obtain a specific combination of trading rules that could be employed to achieve the best market performance no matter what the objective function is.

For the multi-objective optimization, we find that the results are much similar to those obtained in the single-objective optimization, such as the large tick sizes, the connection between tick sizes and multi-price rules, and the most often obtained time priority rule. Similar to the single-objective optimization results, the best-found trading rules differ significantly with various pairs of objective functions. There is, however, a difference that the market transparency shows a clear trend in the multi-objective optimization, while it only presents ambiguous results in single-objective optimization experiments.

7.2 Policy implication

As we only concentrate on how the market structure would impact market performance, and the primary cause of the market efficiency is the market discipline imposed on traders rather than traders behavior, we employ near-zero-intelligence traders. Though traders in financial markets may not be the same as the near-zero-intelligence traders, the model developed in this research could closely mimic the real markets and generate realistic and efficient results. Based on our results, we can practically provide some suggestions on market structure to exchanges or policy makers.

From the results of single-objective optimization experiments, a small tick size, which is currently applied by most of the exchanges, should be applied to stimulate trading activities and increase trading volume. However, in order to reduce the bid-ask spread or return volatility, the tick size should be comparatively larger. In addition, a connection between the tick sizes and

multi-price rules has been found. Thus, we suggest exchanges to employ a tick size which suits best to their concern and consider the relation between the tick size and multi-price rule while designing their market structure.

Although there is no particular priority rule and multi-price rule that are desirable for all exchanges with different performance measures, the time priority and the closest multi-price rule are the most frequently obtained rules in the various experiments. The time priority rule, in conjunction with price priority rule, is always enforced in many real markets, however, the closest multi-price rule is rarely seen. Most of the exchanges select the price which minimize the unexecuted volume as the trading price. Therefore, we suggest exchanges to adopt the closest multi-price rule to improve the market performance.

Other trading rules, such as the intervention of market maker and market transparency, show ambiguous results in different experiments. We, therefore, suggest that these rules do not have a significant influence on exchanges' performance.

In order to consider the various concerns of market participants, exchanges should also employ multiple objective functions when designing their market structure. According to our results, most of the best-found trading rules obtained in the multi-objective optimization experiments are much consistent with those attained in the single-objective optimization. For instance, a large tick size is still recommended; the relation between the tick size and multi-price rule remains unchanged; and the time priority is, again, found to be the most often attained priority rule. However, different from the ambiguous results discussed above, we recommend exchanges to have a small faction of informed traders in order to achieve the best performance when multiple objectives are used.

These are the recommendations for exchanges to improve their market structures based on our experimental results. However, as the best-found combination of trading rules differ significantly with various performance measures,

we suggest exchanges to investigate the best market structures based on their own concerns using our methodology and results.

7.3 Research contributions

From an academic perspective the research does contribute to the growing literature on agent-based modelling related to mechanism design. Interest in mechanism design has been sparked by the recent auction of 3G telecommunications licences, but also applications to procurement processes in the public sector and the greater involvement of private companies in the operation of public services, e.g. railways, schools or computing facilities, require efficient mechanisms to ensure efficient outcomes.

The current market microstructure theory does mostly not offer direct advice on the optimal market structure and where it can be inferred from other results. Given the often complex interactions between trading rules as well as the strong behavioral assumptions made in these models, it is difficult to determine the optimal market design from these results. In addition, the traditional approach of mechanism design seeks closed-form solutions to evaluate the performance of economic institutions. It fails to measure the dynamic behavior in the real-world markets. In addition, most of the current literature on mechanism design uses game theoretic concepts to compare market structures, which is appropriate in cases where only a few market participants interact. However, the financial markets contain a very large number of traders buying and selling, which makes the situation much more complicated. Therefore, mechanism design which uses game theory to evaluate market structures is in most cases not able to address the complex problems faces by exchanges. Moreover, the research on mechanism design usually addresses a very specific problem, no general framework exists, and comparisons between two or more market structures are limited. Furthermore, the current agent-based models used in mechanism design studies only focus on a

single or a small number of related trading rules on specific aspects of market performance given a very narrowly defined model.

To the best of our knowledge this is the first study that complements and combines market microstructure theory and mechanism design using agent-based model. The model developed in this study overcomes the above limitations of current models in mechanism design though the use of near-zero-intelligence traders and evolutionary optimization algorithm. This research provides a first step towards determining optimal market structure by using a wide range of trading rules simultaneously, including the tick size, degree of intervention by market markets, priority rules, multi-price rules, and market transparency, presenting a picture of the rules exchange markets should be designed to in a sufficient number of realistic circumstances. Another major advantage of using our approach, besides its ability to evaluate different market structure, is that it allows to modify the traders themselves and let a heterogeneous population of traders interact, which in the case of using analytical models as in the market microstructure literature is only difficult to achieve, if at all. By changing these characteristics we can build markets that are much closer to reality than is usually possible, allowing us to evaluate the market structures under much more realistic conditions that an analytical model would enable us to do. Thus this research would provide a first insight into this important field, hopefully triggering more interest and establishing a wider research agenda.

This research is also relevant from an applied point of view as well as an academic perspective. On the applied side, we note that the market structure of an exchange becomes increasingly important in a more competitive environment. Access by companies to efficient financial markets, of which exchanges are important elements, provides companies access to cheap funds that can be used for further wealth creation and contributes significantly to economic growth in developing as well as developed countries. This research provides exchanges with an insight into how to develop their market structure to become more attractive to investors and companies alike. It attempts to optimize the market structure of financial markets systematically. No other

research agenda does currently offer a systematic approach for exchanges to improve all aspects of their market structure.

Furthermore, while this research only addresses generic scenarios and derives conclusions on the near-optimal market structure from them, the developed framework can be easily extended to allow tailoring the circumstances to the specific needs of an exchange as well as other stake holders and give specific recommendations to improve the market structure to attract more companies and traders. In addition, the results of this research could also be used to derive recommendations to regulators on establishing the desirable market structure, make suggestions to securities issuers on choosing the best exchange for their listing and guide investors on deciding the most suitable exchange for trading.

Therefore, the results we obtain in this project will increase our knowledge of optimal financial market structure, complementing the results obtained in market microstructure theory and mechanism design. In addition, it will be of obvious use for exchanges to improve their market structure in order to enhance their competitiveness and by companies seeking a listing of their securities as well as traders deciding which exchange to use for their investments.

7.4   Limitations and future work

Although this research has successfully explored the evolutionary mechanism design of call auction markets, there are still some limitations. This section discusses the limitations of the research, and provides some suggestions for future study.

First of all, due to the limitation of time and computing resources, we only repeat each computer experiment 20 times. With a larger sample size, we might possibly have slightly different findings with more data, and the Pareto-efficient frontier could be more obvious and contain more points.

In this research, we assume that each trader is allowed to submit orders with a random order size. In future research, it would be worth considering traders submitting different order sizes with considerations given to wealth through trading. In addition, we only have a single type of asset, which could be extended to assets with different risks in future. Some other assumptions used in this work could also be relaxed in future work. For instance, short sales could be allowed, and a trader can be both a buyer and a seller simultaneously, i.e., to buy and to sell items at the same time. In addition, the agents could employ some learning techniques, for example, by considering some form of communication across agents, or transferring information across individuals from generation to generation.

Except for the trading rules we have applied, other trading rules, such as the market segmentation, size of block trades, trading frequency etc. might be employed in future studies. Some other factors affecting the performance of an exchange can also be employed, for example, price informativeness, broker commissions, recurrent reporting costs, transaction taxes, clearing fees, and listing costs.

In future research the proposed framework can easily be extended to include other objective functions, like maximizing the trading profits to traders. Research with more objective functions would balance a wider range of interest in the market and investigate the sensitivity of the optimal trading rules to the different preferences of decision-makers, thereby giving a more complete picture of the influences on market performance. In addition, some traders' behavior, such as investor preferences and psychology, could also be employed into our generalized framework to analyze the impact of traders' behavior on market performance in future research.

Appendix

A. Optimization of market structure with single objective function

We have discussed the best-found market structure for maximizing the trading volume in section 5.2.1. In this section, we will analyze the best-found set of trading rules with other objective functions, including bid-ask spread, trading price and volatility, with the same analysis procedures used in section 5.2.1.

A.1   Minimization of bid-ask spread

The second computer experiment of the market structure optimization is conducted using the quoted bid-ask spread, also named as inside spread, as performance measure. It computes the amount by which the best ask price exceeds the best bid price, which is the difference between the highest price that a buyer is willing to pay and the lowest price that a seller is willing to sell an asset. As the most common measure of trading costs and market quality in market microstructure theory [1], bid-ask spreads can be used to assess the market liquidity and trading costs for investors. Table A.1 shows the resulting value of the spread and trading rules in 20 simulations.

---

[1] See O'Hara (1995)

Simulation	Spread	Fraction specialist (θ)	Fraction revised (α)	Price adjustment (λ)	Fraction informed (γ)	Tick size (t)	Multi-price rule	Priority rule
1	6.407	0.000	0.951	0.797	0.154	5.151	Min imbalance	Time
2	6.447	0.000	0.656	0.298	0.132	4.329	Random	Time
3	6.490	0.003	0.983	0.654	0.176	2.400	Highest	Size
4	6.521	0.003	0.059	0.360	0.029	5.468	Min imbalance	Random
5	6.523	0.002	0.839	0.775	0.067	5.399	Min imbalance	Reverse time
6	6.551	0.002	0.464	0.395	0.095	4.622	Random	Random
7	6.560	0.010	0.310	0.283	0.093	3.533	Lowest	Time
8	6.561	0.002	0.861	0.957	0.153	4.823	Random	Size
9	6.573	0.003	0.613	0.001	0.124	3.297	Lowest	Size
10	6.639	0.002	0.887	0.361	0.181	4.404	Random	Reverse time
11	6.645	0.003	0.849	0.596	0.107	4.045	Lowest	Size
12	6.712	0.002	0.972	0.756	0.041	3.080	Lowest	Random
13	6.718	0.050	0.463	0.881	0.120	3.874	Lowest	Random
14	6.735	0.045	0.431	0.651	0.120	4.150	Lowest	Size
15	6.777	0.002	0.891	0.665	0.105	4.421	Random	Reverse time
16	6.793	0.011	0.947	0.090	0.131	3.787	Lowest	Random
17	6.927	0.080	0.747	0.342	0.171	5.109	Min imbalance	Time
18	6.970	0.004	0.455	0.292	0.070	3.876	Lowest	Random
19	7.152	0.141	0.186	0.216	0.577	5.116	Min imbalance	Time
20	7.239	0.008	0.508	0.763	0.122	2.598	Lowest	Time
Mean	6.697	0.019	0.654	0.507	0.138	4.174	Lowest*	Random*
StDev	0.227	0.036	0.280	0.275	0.111	0.889		

Tab. A.1: Values of the minimized bid-ask spread and the best-found trading rules for each optimization simulation * The mode is shown instead of the mean for priority rules and multi-price rules.

A.1.1 Tick size

Tick size has several impacts during trading, see section 2.1.4.1. As it represents the cost of getting inside other competitors quote, the tick size affects the motivation of submitting limit orders. In addition, the tick size has an impact on the spreads. As reported in many empirical investigations, e.g. Goldstein and Kavajecz (2000) and Jones and Lipson (2001), amongst others, the bid-ask spread declined dramatically by about 25% with the reduction of tick size from 1/8 to 1/16 dollar on the New York Stock Exchange (NYSE). However, our simulation results suggest that in order to minimize the bid-ask spread the best-found tick size should be in the range from 2.4 to 5.469. This is significantly higher than the current tick adopted in the markets. Therefore, instead of reducing tick size on the stock exchange, it is suggested to be increased to minimize the inside spread. This might be because when tick size is large, quotes of many orders might be centered or close to each other around a few ticks. So, after completion of trading, the bid-ask spread of the remaining orders would be small.
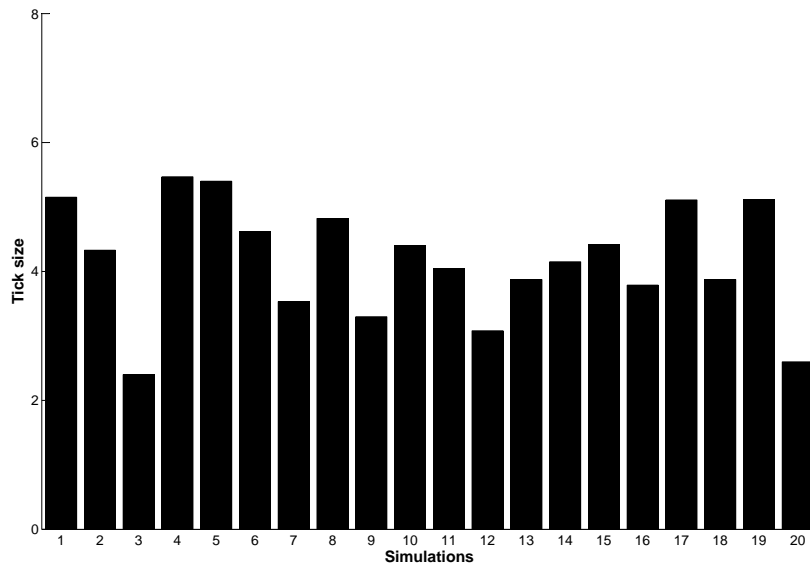
Fig. A.1: the best-found tick size for minimizing bid-ask spread

A.1.2 Priority rules and multi-price rules

From the computer experiments, we do not observe a unique combination of priority rule and multi-price rule that could apply to minimize the bid-ask spread. The most frequently used priority rule is the one to select orders randomly; while the most frequently obtained multi-price rule is the one to choose the lowest price. Similar to the results of maximizing trading volume, the only priority rule that cannot be applied to minimize the spread is the rule of pro-rata. For multi-price rules, three rules could not achieve the desirable structure, including choosing the price both closest and furthest to the previous one, and the price with maximum order imbalance. Like the results of maximizing trading volume, we find a similar pattern between the multi-price rule and tick size. While the optimum tick size is minimum, the multi-price rule of selecting the highest price is applied; when tick size increases, the the multi-price rule of selecting the lowest price is obtained; with even higher the best-found tick sizes, the random selection rule is used, and finally, the multi-price rule of selecting the price with minimum order imbalance is desirable with the largest tick size to minimize the spread.
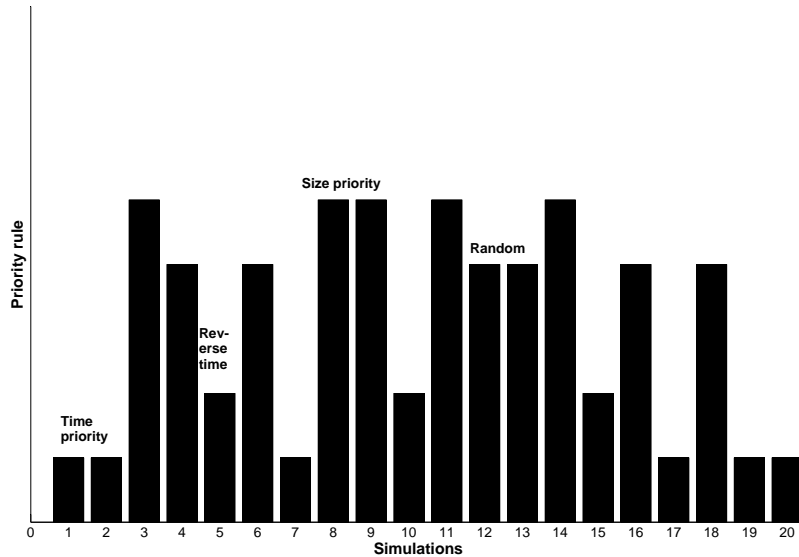
Fig. A.2: the best-found priority rule for minimizing bid-ask spread
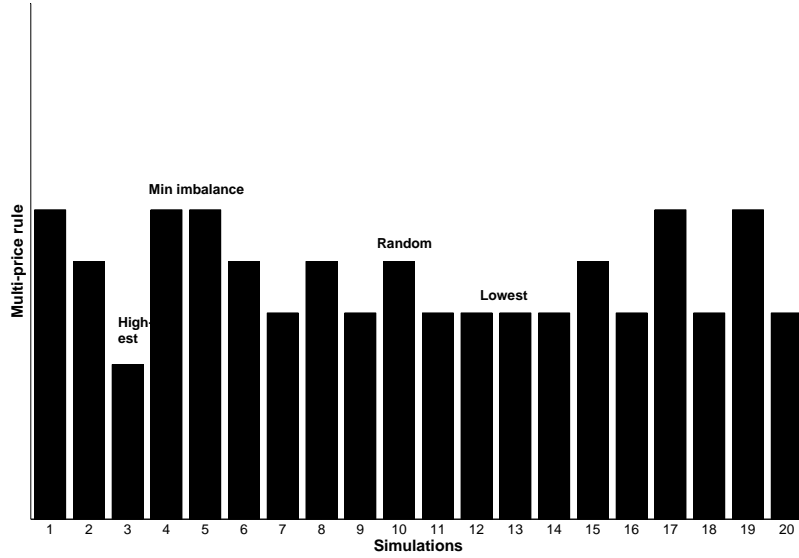
Fig. A.3: the best-found multi-price rule for minimizing bid-ask spread

A.1.3 Transparency

With respect to market transparency, the best-found market structure indicates that approximately 14% traders should be able to access information on the order book. All the obtained the best-found fractions of informed traders are below 20%, except for one case having the maximum fraction of informed traders of 57.7%, which is associated with the highest best-found fraction of order imbalance traded by market maker. The informed traders are allowed to revise their own order size with the average adjustment fraction of 65.4% of the order imbalance, suggesting a rather large revision. However, the adjustment fraction varies with a quite large dispersion, suggesting that no particular desirable value exists. This result implies that information on the order book should not be disseminated too widely in order to ensure a small bid-ask spread. Moreover, we observe that while the fraction of informed traders is minimum, the fraction of order size revised by the informed traders is the smallest as well. However, this relation does not hold when the fraction of informed traders is maximum. Empirically, Flood and Mahieu (1999) discover that the pre-trade transparency can narrow spreads. How-

ever, one can argue that transparency can make it difficult to supply liquidity, because traders may be reluctant to submit limit orders, as the disclosure may convey information in the market which makes the price move against the trader's position.

A.1.4 The intervention of market maker

According to the simulation results, a small fraction of only 1.9%, on average, of order imbalance should be traded by market maker, with a low dispersion. The parameter value is found to be associated with the fraction of informed traders: the more the informed traders, the higher fraction of imbalance traded by market maker, e.g., with the maximum fraction of informed traders, the highest fraction of order imbalance should be traded by market maker. Similar to the results of maximizing volume, it is hard to obtain a particular optimum value of price adjustment of market maker to minimize bid-ask spread. The price adjustment factor λ widely spreads between 0.001 and 0.957. Thus, we could not determine the importance of inventory held by the market maker for the consideration of his price setting based on the results.

A.1.5 Summary of results

To sum up, in order to obtain the minimal inside spread, the stock exchange should allow only a small proportion (about 14%) of traders access to information in the order book, and permit a low fraction (about 1.9%) of order imbalance traded by market maker. The use of priority rules and multi-price rules, however, does not show any relationship to the inside spread. Tick size is also found to be ambiguous in the best-found market structure. We can thus conclude that only the fraction of informed traders and the extend of market making seems to have significant impacts on bid-ask spread.
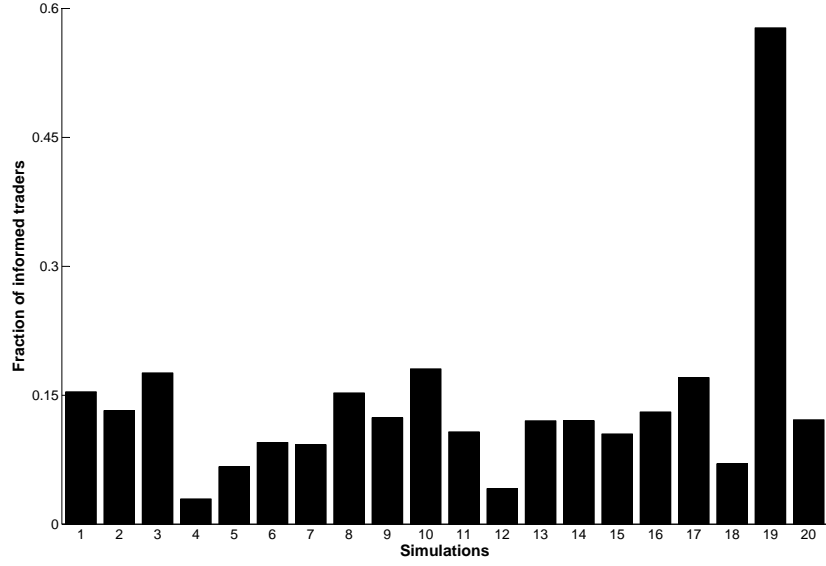
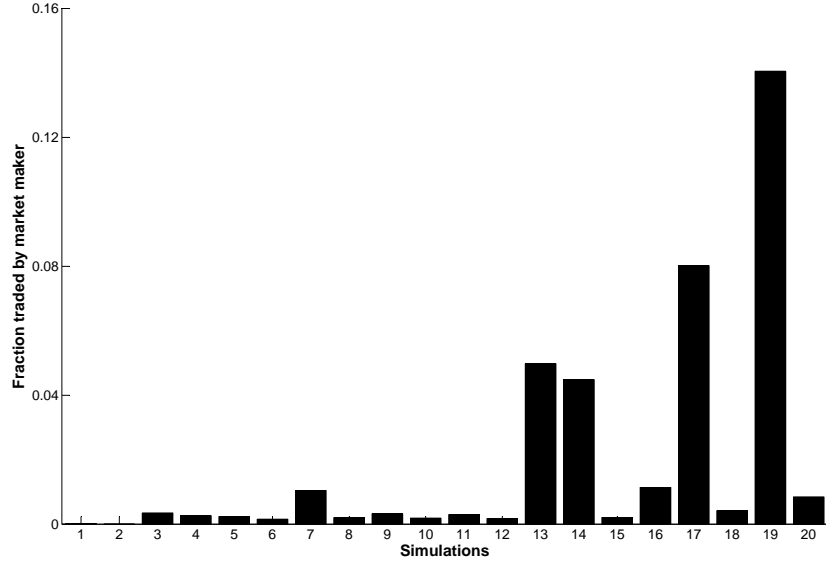(a) the best-found fraction of informed traders
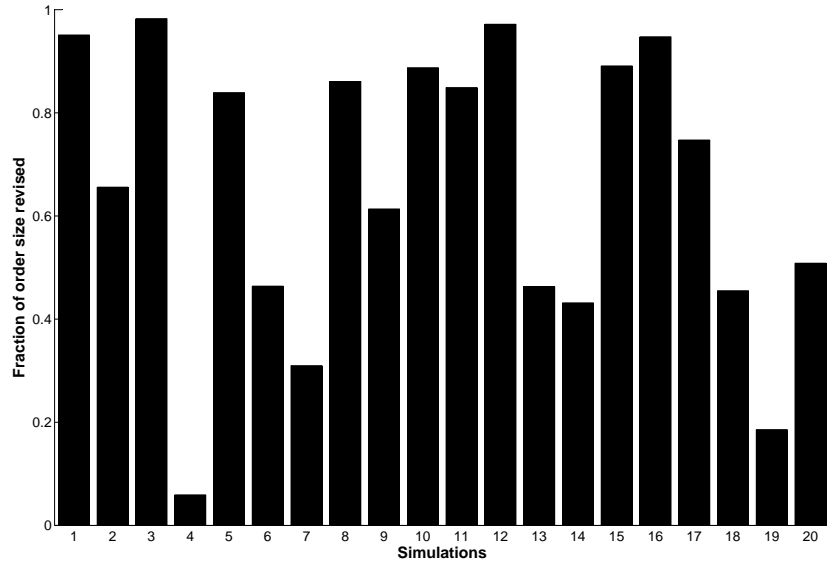
(b) the best-found fraction of order size revised

Fig. A.4: Simulation results of market transparency for minimizing bid-ask spread

A.2 Maximization of trading price

Trading price, one of the most important factors in trading activities, is another objective function used to study the optimal market structure. A

(a) the best-found fraction of order imbalance traded by market maker
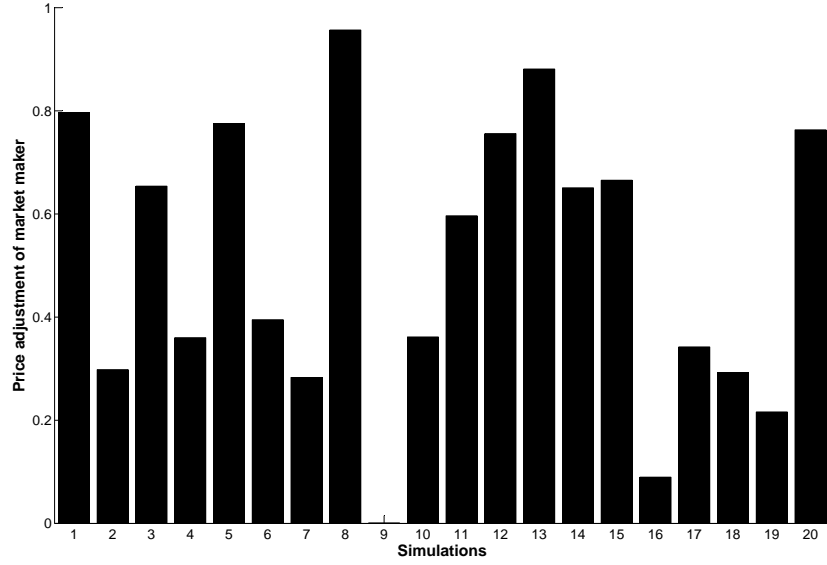
(b) the best-found price adjustment of market maker

Fig. A.5: Simulation results of intervention of market maker for minimizing bid-ask spread

market trader would realize more profit in markets with increasing prices. Therefore, we intend to obtain the best combination of trading rules to maximize the trading price. Table A.2 shows the the best-found trading rules in

20 experimental simulations.

A.2.1 Tick size

Table A.2 shows that the best-found tick sizes from the simulation fluctuate between 0.057 and 1.623. Most of the tick sizes distribute in the region of 0 to 0.3, while others are ranging from 0.3 to 1.623, see figure A.6. In accordance with the results of maximizing volume and minimizing spread, the best-found tick sizes for maximizing the price are correlated with the selection of the multi-price rule. If the tick size is smaller than one, multi-price rules of choosing the closest price to the previous price is applied; while the multi-price rule of choosing the furthest price to the previous price is obtained if tick size is larger than one. However, we do not observe a clear trend of the trading price variation with different tick sizes.
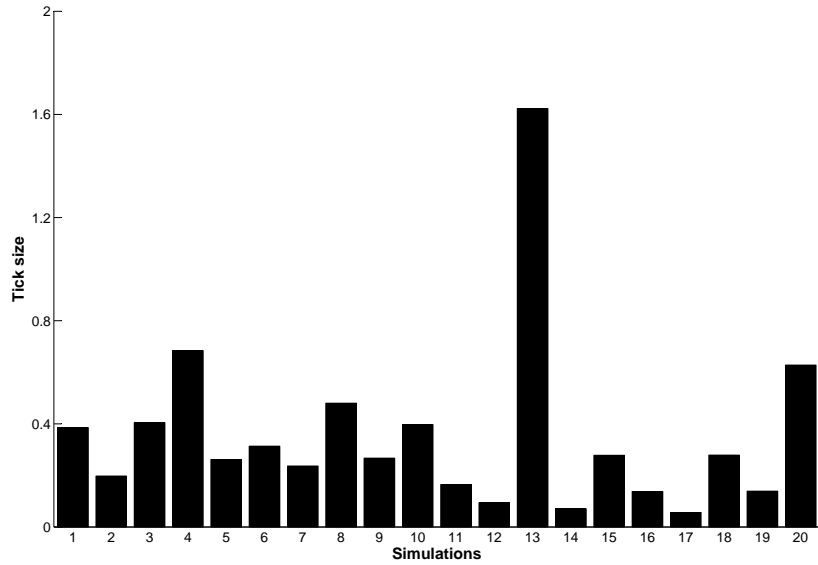

Fig. A.6: the best-found tick size for maximizing trading price

Simulation	Price	Fraction specialist (θ)	Fraction revised (α)	Price adjustment (λ)	Fraction informed (γ)	Tick size (t)	Multi-price rule	Priority rule
1	148.530	0.000	0.806	0.886	0.388	0.385	Closest	Reverse time
2	149.264	0.005	0.248	0.471	0.812	0.198	Closest	Time
3	150.515	0.002	0.012	0.062	0.103	0.404	Closest	Reverse time
4	152.492	0.008	0.813	0.386	0.359	0.684	Closest	Random
5	158.188	0.008	0.272	0.846	0.410	0.262	Closest	Time
6	159.589	0.006	0.052	0.094	0.875	0.313	Closest	Reverse time
7	160.477	0.003	0.762	0.863	0.545	0.236	Closest	Time
8	163.608	0.006	0.262	0.952	0.884	0.480	Closest	Reverse time
9	165.627	0.221	0.705	0.513	0.942	0.267	Closest	Time
10	166.045	0.005	0.778	0.355	0.577	0.397	Closest	Reverse time
11	166.333	0.005	0.444	0.960	0.737	0.164	Closest	Time
12	170.784	0.005	0.847	0.511	0.743	0.094	Closest	Time
13	171.157	0.003	0.048	0.938	0.161	1.623	Furthest	Reverse time
14	171.976	0.003	0.537	0.952	0.306	0.071	Closest	Time
15	173.839	0.269	0.580	0.042	0.966	0.278	Closest	Time
16	175.534	0.008	0.470	0.206	0.790	0.138	Closest	Time
17	177.194	0.370	0.606	0.405	0.659	0.057	Closest	Time
18	178.214	0.253	0.148	0.080	0.249	0.279	Closest	Time
19	178.483	0.006	0.568	0.873	0.970	0.139	Closest	Time
20	180.018	0.770	0.233	0.217	0.608	0.628	Closest	Random
Mean	165.893	0.098	0.460	0.531	0.604	0.355	Closest*	Time*
StDev	14.574	0.195	0.282	0.347	0.277	0.344		

Tab. A.2: Values of the maximized trading price and the best-found trading rules for each optimization simulation

* The mode is shown instead of the mean for priority rules and multi-price rules.

A.2.2 Priority rules and multi-price rules

We do not obtain a specific combination of the best-found priority rule and multi-price rule. Almost all of the obtained multi-price rules are to choose the price closest to the previous price. In one particular case mentioned above that when the best-found tick size is larger than 1, the multi-price rule of selecting the price furthest from the previous price is applied. For the case of imbalance between buy and sell orders at the transaction price, there are three desirable rules could be used to maximizing the trading price, namely the time priority, reverse time priority, and random selection, of which the time priority rule in which orders submitted earlier will receive priority to be filled is the most frequently obtained.
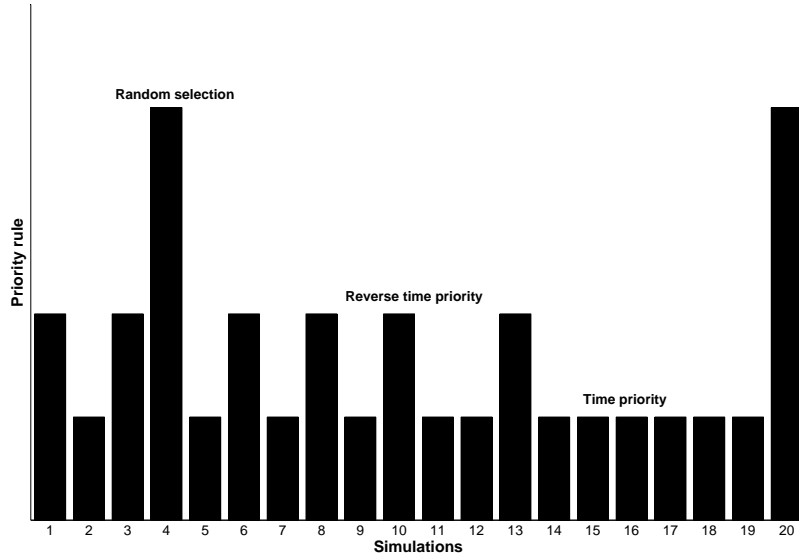
Fig. A.7: the best-found priority rule for maximizing trading price

A.2.3 Transparency

Different from the above results where a small fraction of traders should access information of the order book to maximize volume or minimize spread,
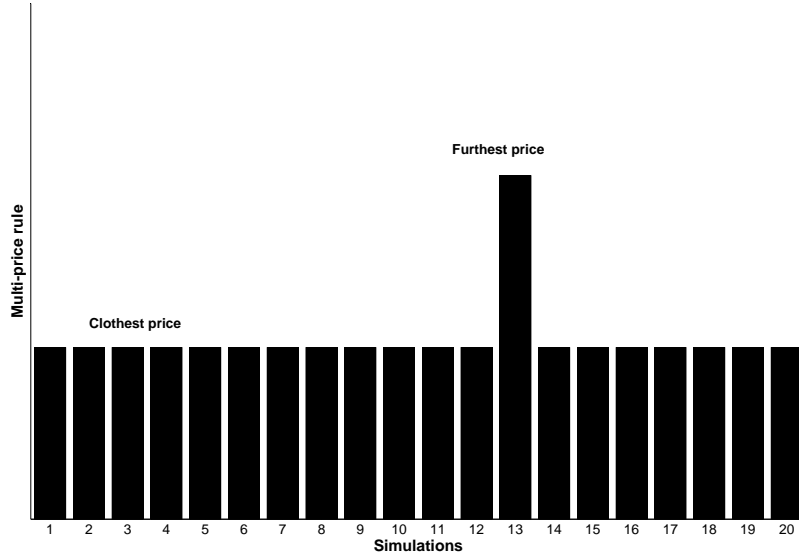
**Fig. A.8:** the best-found multi-price rule for maximizing trading price

the fraction of informed traders shows an ambiguous impact on trading prices. On average, 60.4% of traders are suggested to be informed, however, this fraction varies in a wide range from 10.3% to 97%. Moreover, the fraction of order size revised by informed traders is also widespread, showing no clear pattern.

A.2.4 The intervention of market maker

Based on the 20 simulation results, we notice that while the fraction of imbalance traded by market maker is increasing, the trading price is increasing as well for most cases, see A.10(a). For instance, when the smallest fraction of imbalance is traded by market maker is applied, the corresponding price is minimal. On the other hand, while the highest fraction of imbalance is traded by market maker, the price raises to the maximum. However, it is difficult to identify a particular trend or desirable value for the price adjustment of market maker, due to the wide range of fluctuation.
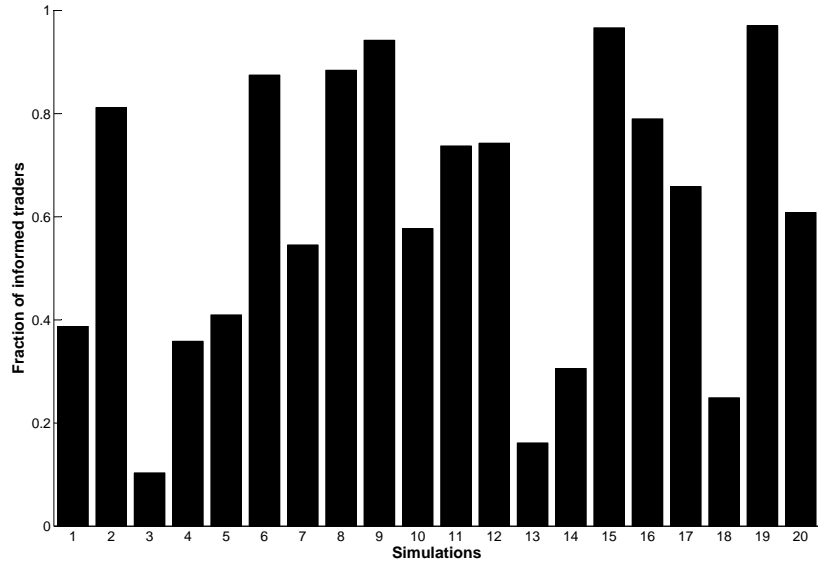
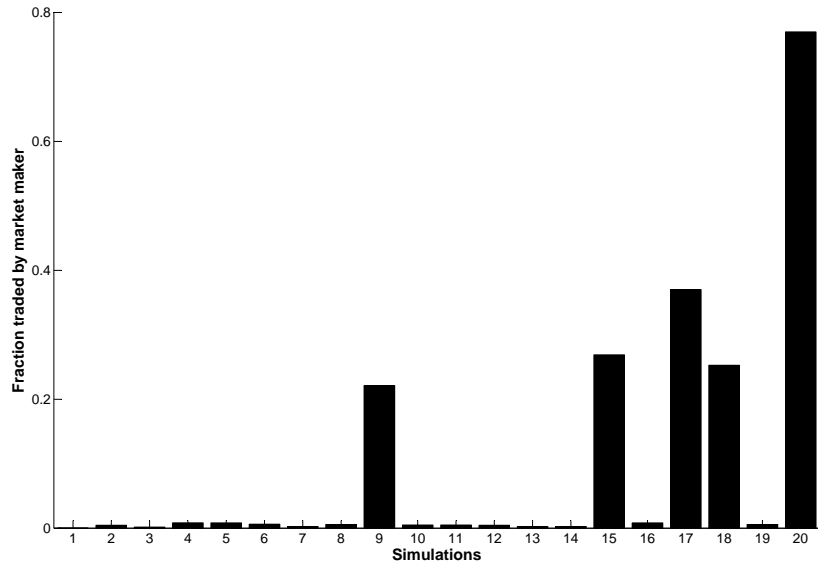(a) the best-found fraction of informed traders
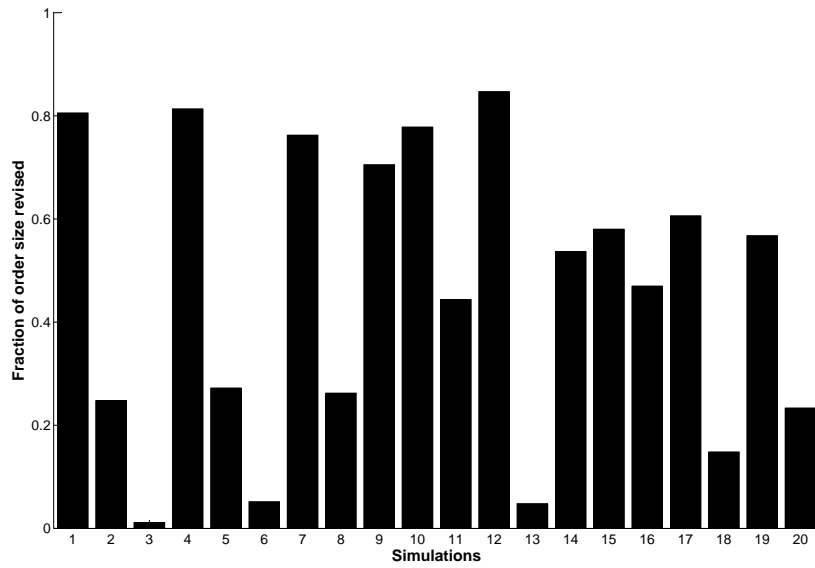

(b) the best-found fraction of order size revised

Fig. A.9: Simulation results of market transparency for maximizing trading price

A.2.5 Summary of results

In summary, we find that a tick size of approximately 0.355 should be employed to achieve the maximal trading price; the selection of tick size is in

(a) the best-found fraction of order imbalance traded by market maker
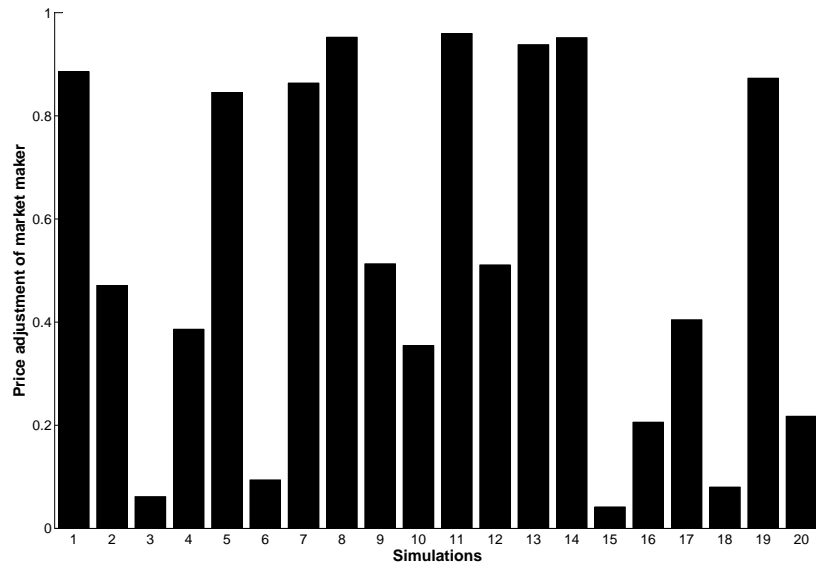
(b) the best-found price adjustment of market maker

Fig. A.10: Simulation results of intervention of market maker for maximizing trading price

connection with the choice of multi-price rule, and the intervention of market making has a positive influence on trading price; the impacts of transparency and priority rule are, however, uncertain to the trading price. Therefore, the

main effect on trading price arises from the tick size, multi-price rule, and intervention of market maker.

A.3   Minimization of return volatility

Volatility, which is frequently referred to as the standard deviation of the returns within a specific time period, is often used to measure the risk of a financial instrument over a period as well as the liquidity of the market. Most of market participants consider the volatility as a negative in that it signifies risk and uncertainty. Therefore, the last objective function we employed for the single objective optimization experiments is to minimize the return volatility, see table A.3 for the results.

A.3.1   Tick size

The best-found tick size is selected from a widespread range (from 0.007 to 3.451), which is correlated with the choice of multi-price rule. Most of the obtained tick sizes, see figure A.11, are quite small. Under these circumstances, choosing the price closest to the previous price is always been applied as the multi-price rule. When the best-found tick size increases to 1.832, the multi-price rule of choosing the price furthest from the previous price is found to minimize the volatility. And, finally, when the best-found tick size reaches its maximum of 3.451 in this simulation, the multi-price rule of choosing the highest price is then selected. Meanwhile, with the highest tick size the fraction of imbalance traded by market maker would achieve the maximal value. However, there is no correlation while tick size is minimum based on the results. Furthermore, we find that the volatilities observed from the simulation experiments are, to a certain extent, related to the tick size. Based on our results, the stock returns are more volatile while tick size increases. As tick size increases, liquidity supply could be reduced, which is associated with higher volatility, since larger tick size would lead to market participants' reluctance

Simulation	Volatility	Fraction specialist (θ)	Fraction revised (α)	Price adjustment (λ)	Fraction informed (γ)	Tick size (t)	Multi-price rule	Priority rule
1	27.700	0.004	0.761	0.798	0.278	0.428	Closest	Reverse time
2	28.248	0.001	0.457	0.798	0.425	0.092	Closest	Time
3	29.414	0.002	0.623	0.802	0.829	0.007	Closest	Time
4	30.027	0.008	0.739	0.689	0.634	0.837	Closest	Random
5	30.533	0.003	0.473	0.851	0.927	1.024	Closest	Size
6	30.697	0.005	0.405	0.023	0.993	0.506	Closest	Reverse time
7	30.854	0.007	0.226	0.608	0.254	0.472	Closest	Reverse time
8	30.994	0.000	0.765	0.662	0.814	0.514	Closest	Reverse time
9	31.164	0.002	0.194	0.570	0.926	0.210	Closest	Time
10	31.465	0.001	0.778	0.134	0.823	0.172	Closest	Time
11	31.874	0.003	0.839	0.777	0.504	0.586	Closest	Reverse time
12	32.100	0.006	0.450	0.100	0.872	0.152	Closest	Time
13	32.154	0.002	0.910	0.990	0.712	0.663	Closest	Random
14	32.544	0.005	0.373	0.351	0.367	0.056	Closest	Time
15	34.155	0.002	0.392	0.747	0.767	0.168	Closest	Time
16	35.264	0.002	0.514	0.688	0.472	0.268	Closest	Time
17	37.895	0.008	0.633	0.269	0.469	0.978	Closest	Size
18	38.111	0.000	0.432	0.285	0.753	0.186	Closest	Time
19	39.547	0.004	0.923	0.662	0.740	1.832	Furthest	Reverse time
20	40.572	0.009	0.360	0.939	0.974	3.451	Highest	Size
Mean	32.767	0.004	0.562	0.587	0.677	0.630	Closest*	Time*
StDev	3.676	0.003	0.221	0.299	0.235	0.794		

Tab. A.3: Values of the minimized return volatility and the best-found trading rules for each optimization simulation

* The mode is shown instead of the mean for priority rules and multi-price rules.

of supplying liquidity. Although the theoretical predictions regarding the influence of tick reduction on volatility are mixed, e.g., Ronen and Weaver (1998) reviews the two different opinions: "The first argument, which predicts an increase in observed volatility is based on the premise the depths will decrease as a result of the tick reduction ... Conversely, theoretical justification exists in support of volatility decreases (following a reduction in tick size)", some research, such as Harris (1990a), Ronen and Weaver (1998), and Bessembinder (2000) confirm that observed volatility decreases with a reduction on tick size, which is consistent with our result.
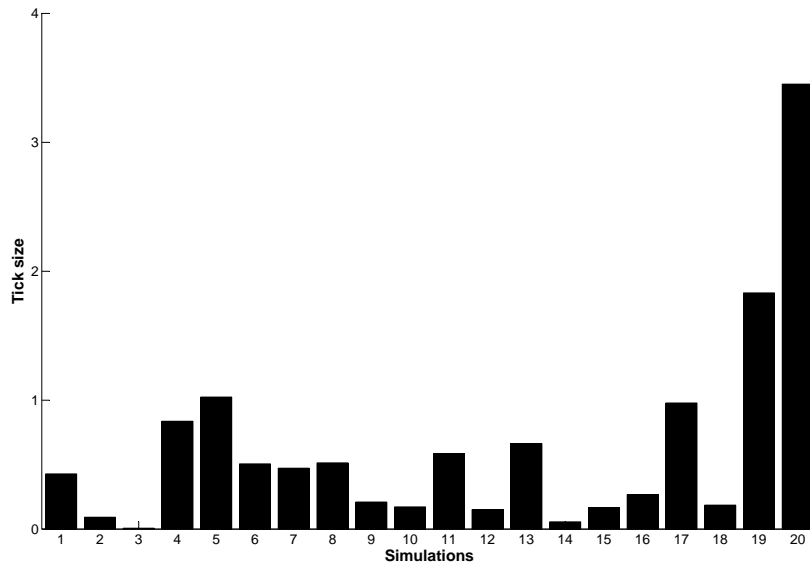
Fig. A.11: Simulation results of tick size for minimizing volatility

A.3.2 Priority rules and multi-price rules

Similar from the results of maximizing trading volume, it is difficult to identify a specific priority rule to achieve the best-found market structure of minimizing volatility. All the various priority rules, exclusive of the pro-rata rule, are obtained to achieve the optimum result, and the most frequently obtained is time priority rule. For the case of multiple prices at which the

trading volume is maximal, the multi-price rules that choosing the price both closest and farthest to the previous price, and selecting the highest price could be used to achieve the optimum market structure in terms of minimized volatility. And, the always observed best-found multi-price rule is selecting the price closest to the previous price. In addition, we find a relationship between the tick size and the multiple price rule, which has been discussed above.
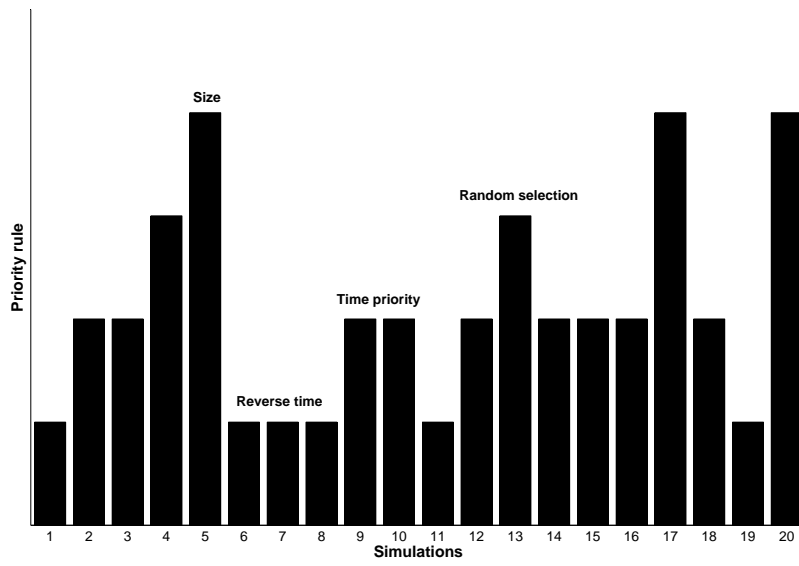

Fig. A.12: the best-found priority rule for minimizing volatility

A.3.3 Transparency

From the computer experiments, we could not find a particular desirable fraction of traders who can access pre-trade information on the order book. With the mean value of 67.7%, the best-found fraction of informed traders ranges from 25.4% to 99.3%. Moreover, these informed traders are allowed to revise their own order size with the adjustment of an average fraction, 56.2% of the order imbalance. This parameter value, with a dispersion of 22.1%, as
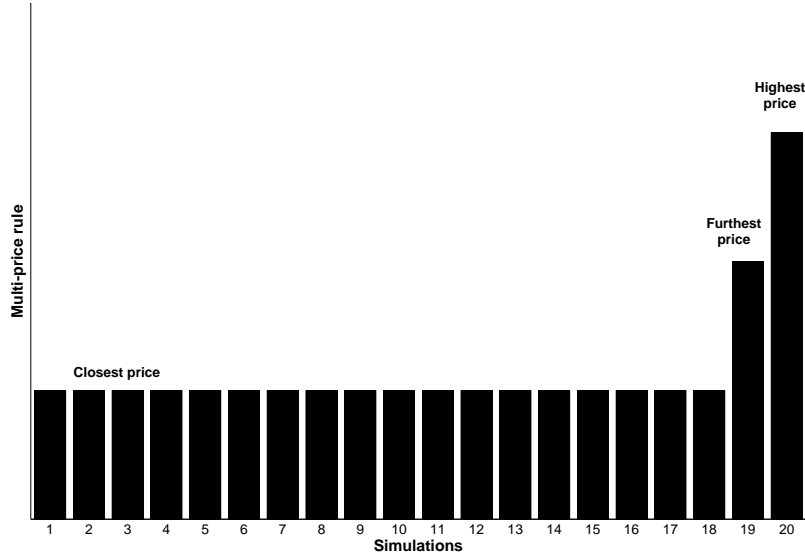
**Fig. A.13:** the best-found multi-price rule for minimizing volatility

shown in table A.3, varies widely from the average value in different simulations. Thus, we could not determine the importance of market transparency for the consideration of volatility minimization.

A.3.4   The intervention of market maker

In respect of the intervention of market maker, the best-found market structure with the objective function of volatility minimization indicates that a very small fraction (0.4% on average) of imbalance should be traded by market maker in order to achieve the optimum results. The resulting parameter values change little from the average. However, the transaction price formation factor of market maker fluctuates from a broad range (between 0.195 and 0.923).
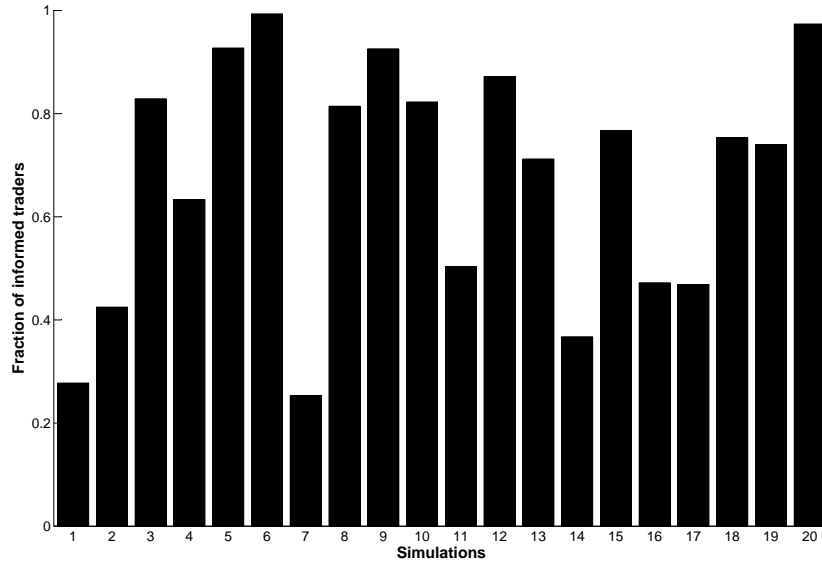
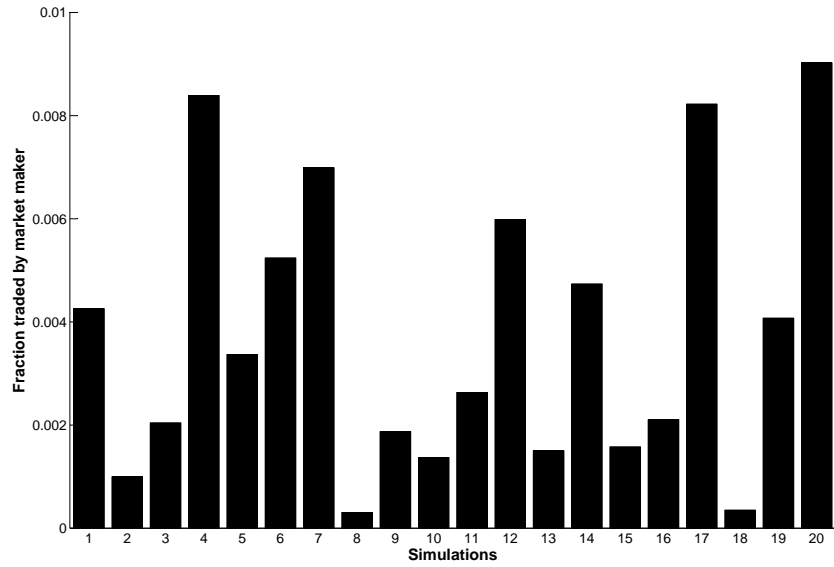(a) the best-found fraction of informed traders
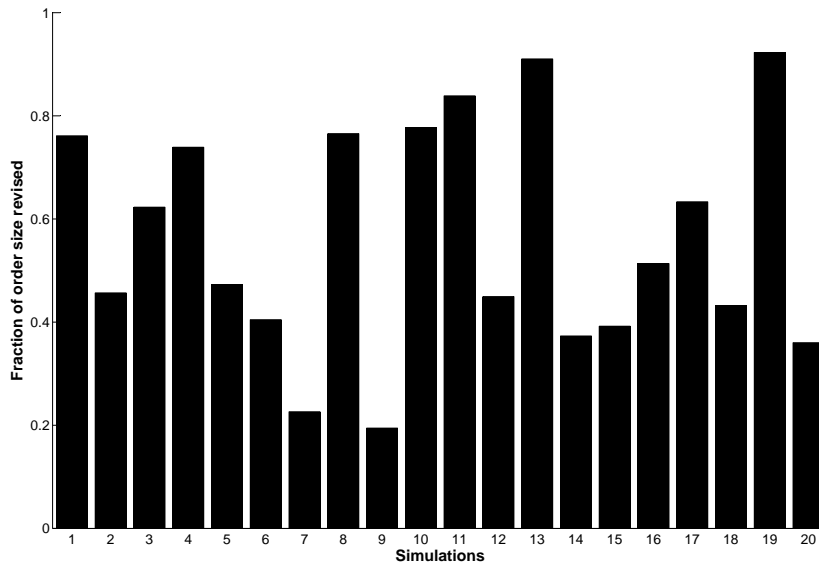

(b) the best-found fraction of order size revised

Fig. A.14: Simulation results of market transparency for minimizing volatility

A.3.5 Summary of results

The main results of this experiment can be concluded as follows. It is difficult to find a clear trend of the best-found fraction of informed traders. On

(a) the best-found fraction of order imbalance traded by market maker for minimizing volatility
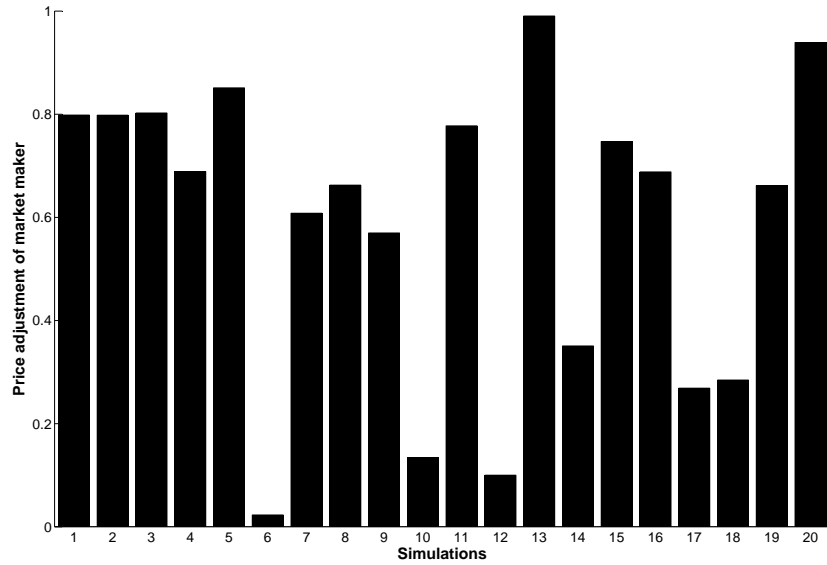
(b) the best-found price adjustment of market maker

Fig. A.15: Simulation results of intervention of market maker for minimizing volatility

the other hand, we observe that a very small proportion (about 0.4%) of imbalance should be traded by market maker. Most of the best-found tick

sizes are small. All the various priority rules, apart from the pro-rata rule, are obtained to achieve the optimum result. And, the selection of multi-price rule is correlated with the value of tick size. In most of the best-found market structures, the multi-price rule of selecting the price closest to the previous price is adopted.

B. Optimization of market structure with multi-objective functions

We have discussed the best-found market structure for maximizing the trading volume and minimizing bid-ask spread in section 6.2.1. In this section, we will analyze the the best-found set of trading rules with other pairs of objective functions, including maximizing trading price and minimizing bid-ask spread; maximizing trading volume and minimizing volatility, and minimizing spread and volatility.

B.1   Trading price and bid-ask spread

Figure B.1 shows the maximized trading price and the minimized bid-ask spread at the final generation for the entire population for all 30 runs of our computer experiments focusing on the region containing the Pareto-efficient frontier. We, once again, notice a trade-off between the maximal trading price and the minimal spread, the approximate location of the Pareto-efficient frontier is sketched by the line to the lower right. This figure clearly shows that a low spread will be associated with a low trading price while a high trading price will impose a large spread. The best combination of trading price and spread will depend on the preferences of the decision-maker and the relative importance of these two aspects. We have identified eight markets that approximately determine the Pareto-efficient frontier, similarly, these points are identified as large points and associated with numbers; the market structures of the eight markets are shown in the following sections.
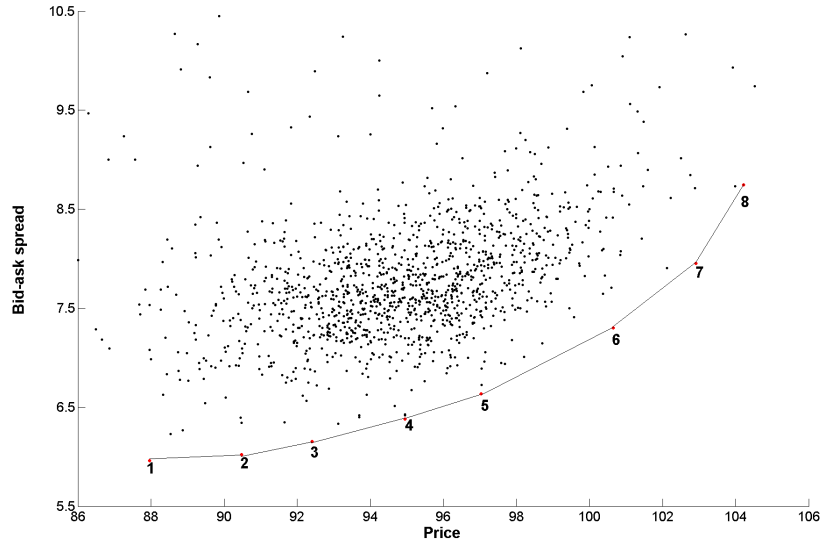
Fig. B.1: Market performance of maximizing trading price and minimizing bid-ask spread after 500 generations

B.1.1   Tick size

The best-found tick sizes seem to be more stable, see figure B.2 floating around the mean value of 3. Compare to the results of single-objective optimization with fitness function of maximizing price, and minimizing spread respectively, the mean value in this experiment is between the average of the best-found tick sizes from the two single-objective optimization experiments (0.355 and 4.174, respectively).

B.1.2   Priority rules and multi-price rules

For all the eight cases the best-found multi-price rule is to select the lowest price, see figure B.3. In the single-objective optimization simulations, the lowest multi-price rule has been obtained in the experiment with fitness function of minimizing bid-ask spread, but never attained in the experiment of maximizing trading price. There are three obtained best-found priority rules, see figure B.4 which are reverse time priority rule, size priority rule,
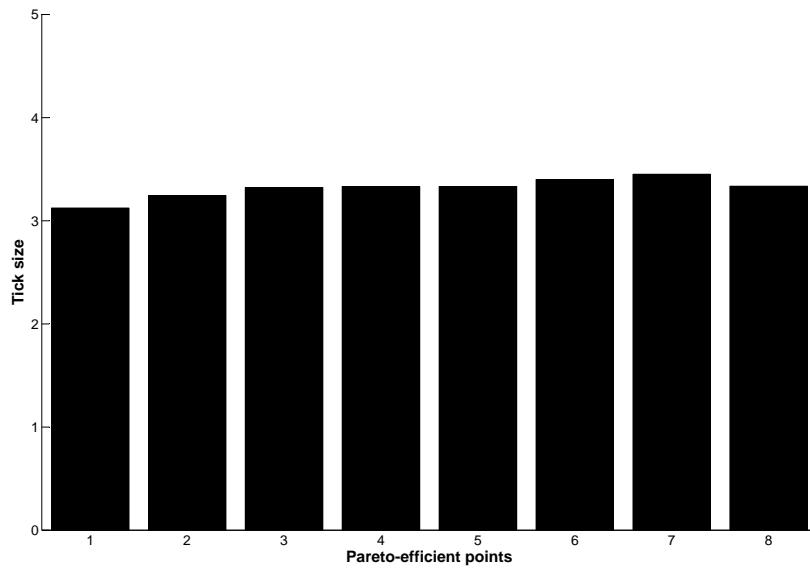
Fig. B.2: the best-found tick sizes for maximizing trading price and minimizing spread
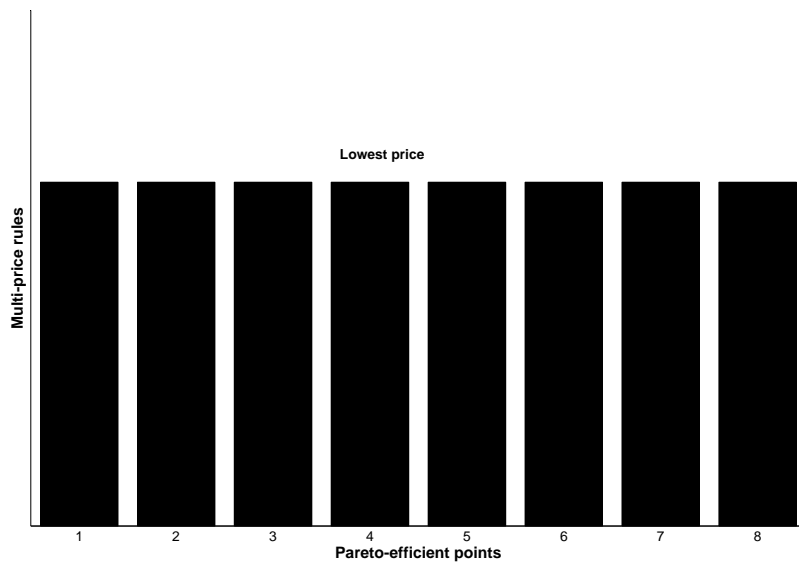

Fig. B.3: the best-found multi-price rules for maximizing trading price and minimizing spread

and time priority rule. Another priority rule, which is the random selection, has also been observed in the two single-objective optimization experiments.
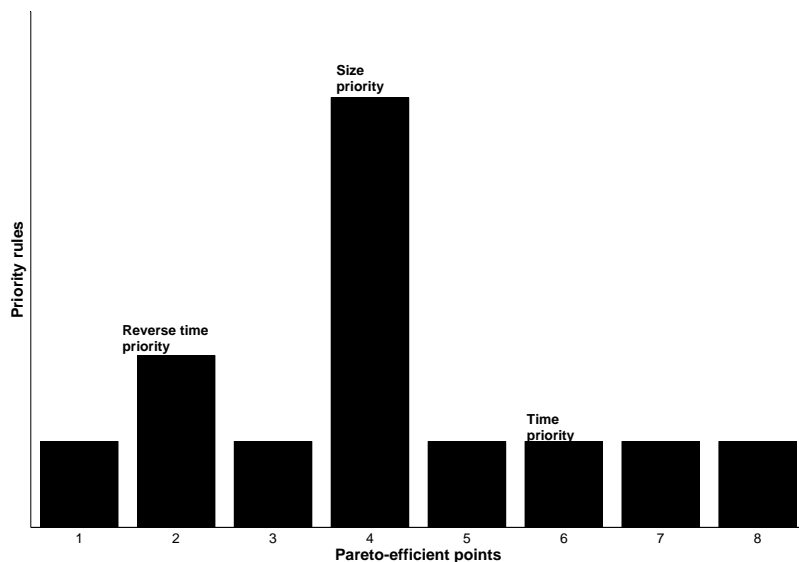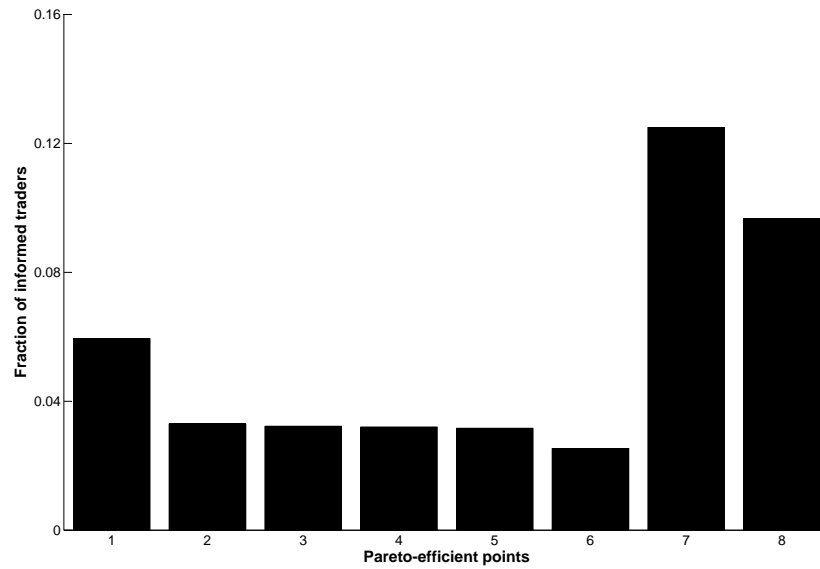
Fig. B.4: the best-found priority rules for maximizing trading price and minimizing spread
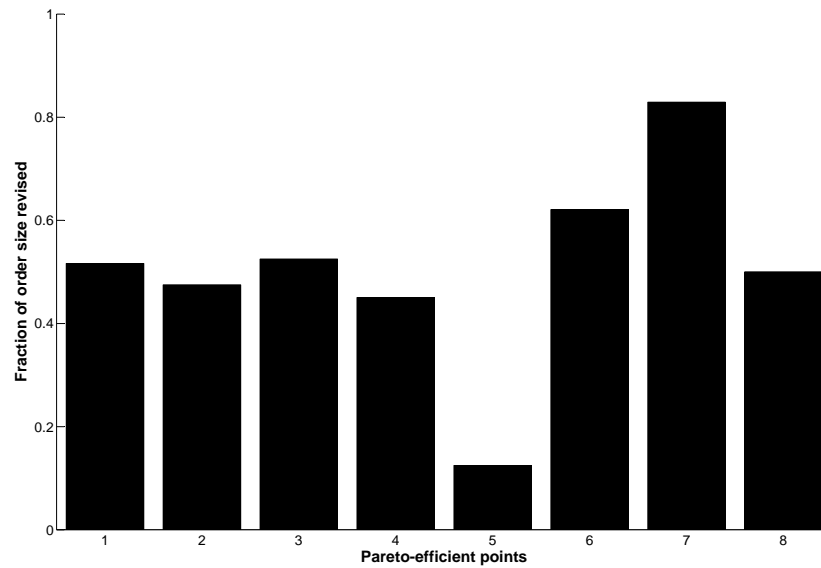
B.1.3 Transparency

The best-found fraction of informed traders shows ambiguous results. It is unable to identify a clear trend in the parameter values as we move along the efficient frontier. The best-found fraction of informed traders is generally small, varying from 2.54% to 12.5%, but comparatively higher than the range obtained in the two single-objective optimization experiments, see figure B.1.3.

B.1.4 The intervention of market maker

Similarly, the best-found market structure with respect to the intervention of market makers is also ambiguous. The degree of intervention by the market
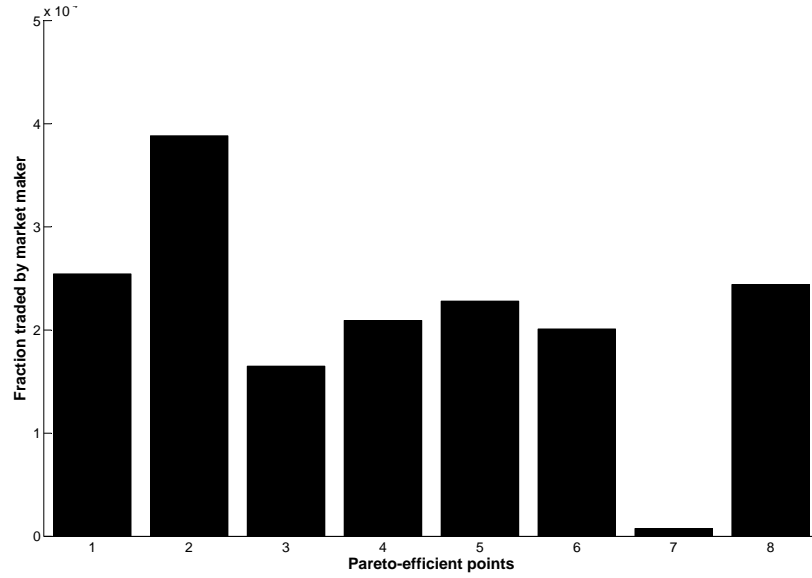
(a) the best-found fraction of informed traders for maximizing trading price and minimizing spread

(b) the best-found fraction of order size revised for maximizing trading price and minimizing spread

Fig. B.5: Simulation results of market transparency for maximizing trading price and minimizing spread

maker is very small, since the maximum fraction of order imbalance traded by market maker is only 0.048%, see figure B.1.4.

(a) the best-found fraction of order imbalance traded by market maker
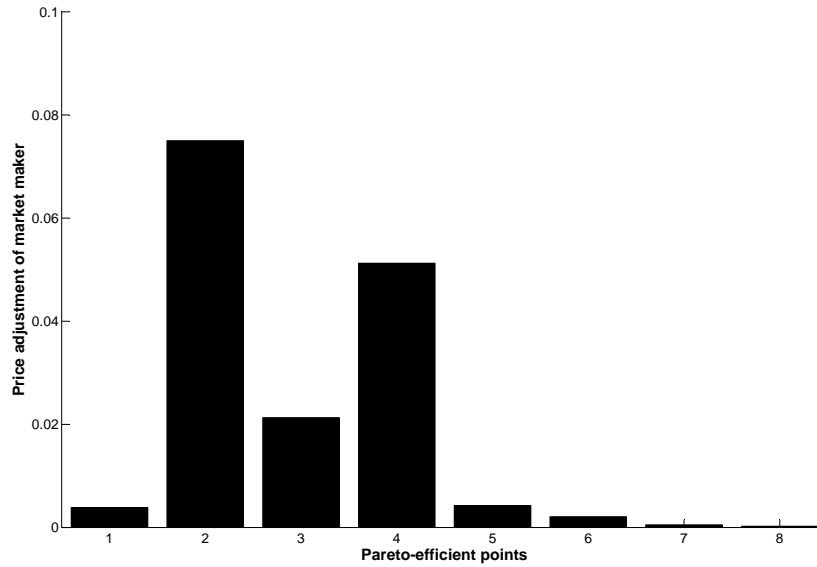
(b) the best-found price adjustment of market maker

Fig. B.6: Simulation results of intervention of market maker for maximizing trading price and minimizing spread

Fig. B.7: Market performance of maximizing trading volume and minimizing volatility after 500 generations

B.1.5 Summary of results

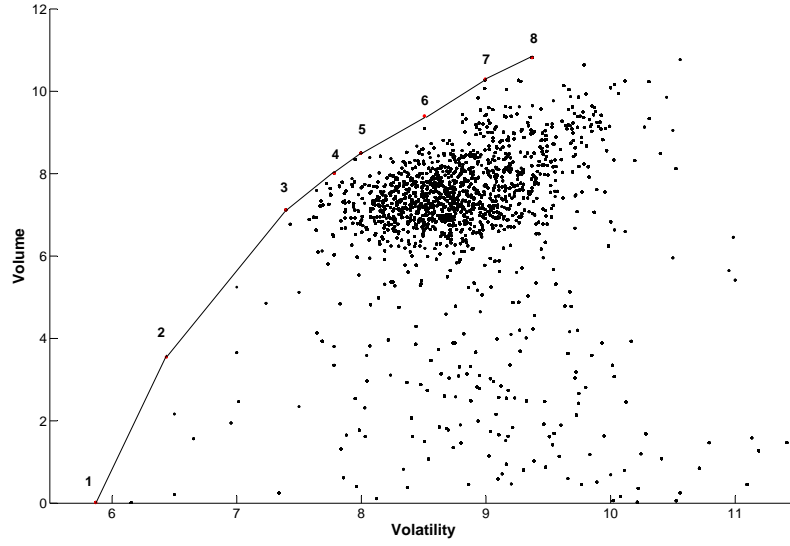We can summarize our results in stating that to attain the best market structure with the objective functions of maximize trading price and minimize bid-ask spread, both extent of intervention of market makers and degree of market transparency are suggested to be low; in all cases the market should employ the multi-price rule of selecting the lowest price, and the tick size should be around 3. In the single-objective optimization experiment of minimizing inside spread, we notice that the most obtained multi-price rule is to select the lowest price, and the extent of market maker intervention is generally low, which are much consistent with our results here, but others such as tick size and priority rules are showing ambiguous results.

B.2 Trading volume and return volatility

We further study the best-found market structure by employing the fitness function of maximizing trading volume and minimizing return volatility. Fig-

ure B.7 displays the best-found trading rules that are evolved after 500 generations in 30 simulations restricted to the area close to the Pareto-efficient frontier, which is plotted by the line. The trade-off relationship between the maximal trading volume and the minimal volatility has been clearly displayed on the figure: high trading volume is associated with high volatility; while low trading volume, on the other hand, is associated with low return volatility.

Similar to the previous analysis, we identified eight markets from the Pareto-efficient frontier, highlighted by large points associated with numbers. The following sections present the selection of trading rules of the eight best-found market structure.

B.2.1   Tick size

Except for the two market structures with high trading volume and large volatility, the best-found tick size is usually very large (ranging from 2.5 to 5), see B.8. Again, consistent with results of the single-objective optimization experiments, the selection of tick size is related to the use of multi-price rules. For instance, when the tick size is very low, the multi-price rule of selecting the nearest price is obtained; while the tick size is very large, the multi-price of selecting the lowest price is employed.

B.2.2   Priority rules and multi-price rules

Different from results of the two respective single-objective optimization experiments (of maximizing trading volume and minimizing volatility), where all priority rules, except for the random selection rule, can be applied to achieve the optimum, in this experiment, we only observe a single best-found priority rule, which is time priority rule, used in the eight best-found market structures, see B.10. Nevertheless, we obtain three best-found multi-price
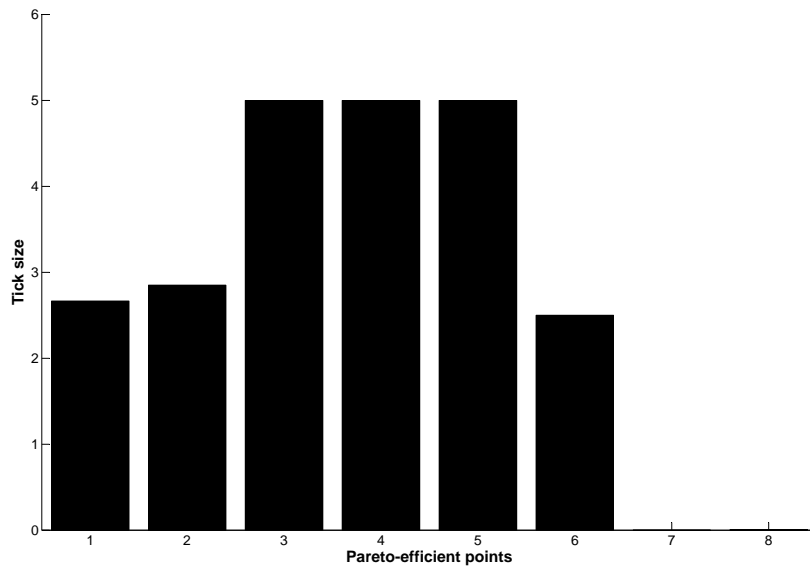
Fig. B.8: the best-found tick sizes for maximizing trading volume and min-
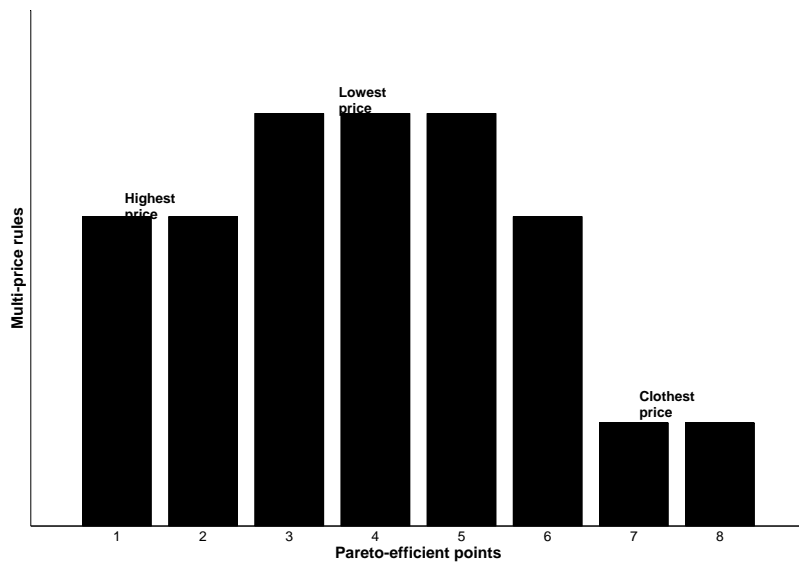imizing volatility

Fig. B.9: the best-found multi-price rules for maximizing trading volume
and minimizing volatility

rules, including choosing the lowest price, the highest price, and the price
closest to the previous price, see B.9
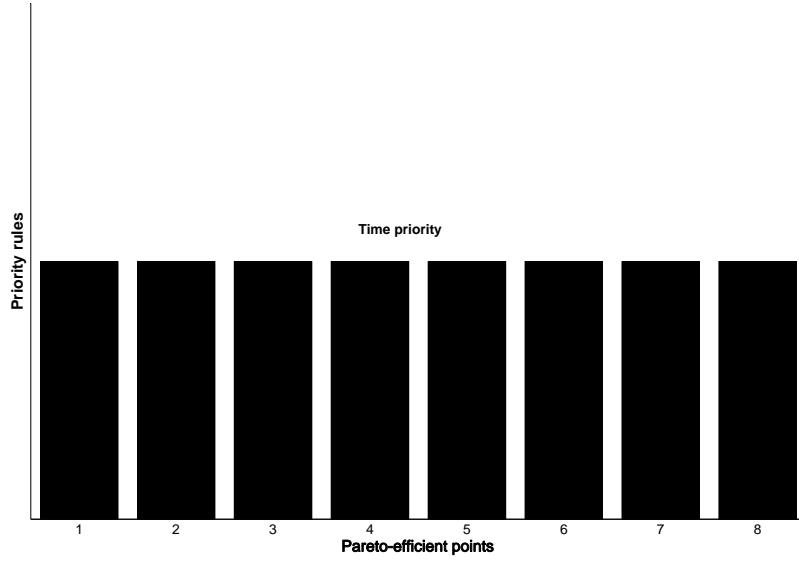
Fig. B.10: the best-found priority rules for maximizing trading volume and minimizing volatility

B.2.3 Transparency

Results of the fraction of informed traders indicate that less traders should be able to access information on the order book to optimizing the trading volume, while more traders should be informed if the decision maker puts great emphasis on small volatility. Overall, the average of the best-found fraction of informed traders is about 2%. As shown in figure B.11(a), when trading volume is the main concern, the fraction of informed traders is comparatively lower. This is consistent with the results of single-objective optimization experiment of maximizing trading volume. However, the fraction of order size used to revise the order by those informed traders seems to be independent of the location on the Pareto-efficient frontier, see B.11(b).

B.2.4 The intervention of market maker

The results shown in figure B.12(a) indicate that the degree of the market maker's intervention should be low, in particular, the market maker should
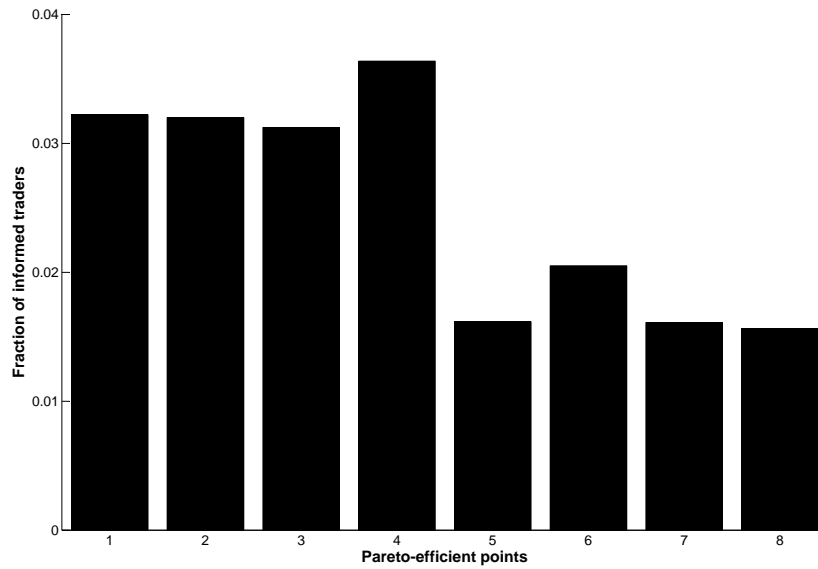
(a) the best-found fraction of informed traders for maximizing trading volume and minimizing volatility
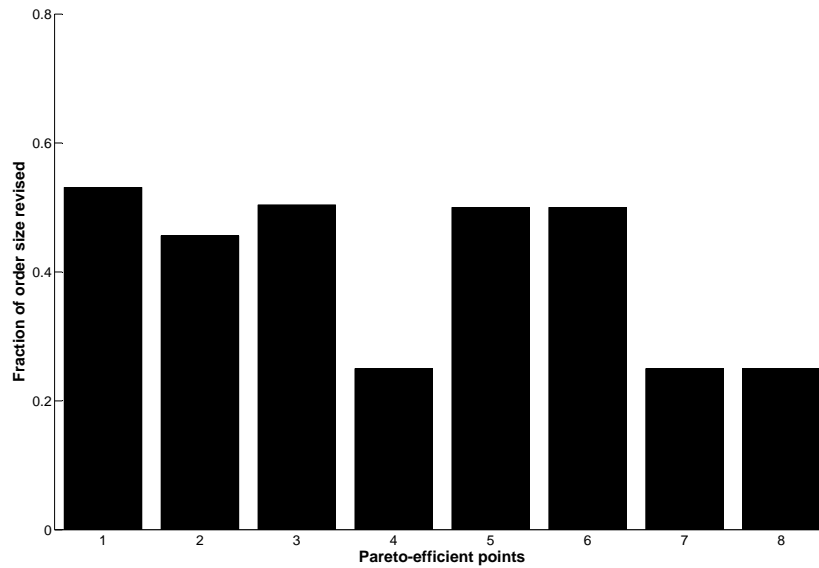
(b) the best-found fraction of order size revised for maximizing trading volume and minimizing volatility

Fig. B.11: Simulation results of market transparency for maximizing trading volume and minimizing volatility

trade a very small portion of the order imbalance. We find that the fraction
of order imbalance traded by market maker should be comparatively higher
to stimulate the trading activities if the main concern of the decision maker
is trading volume. If the main intention is to reduce the return volatility,
the fraction traded by market maker is then comparatively lower. This re-
sult is in consistent with the results of two single-objective experiments of
maximizing trading volume, and minimizing return volatility. However, the
price adjustment of market maker, shown in figure B.12(b), shows no clear
pattern for different market structures.

B.2.5   Summary of results

In summary, in order to achieve higher trading volumes (at the cost of higher
volatility), a smaller tick size and fraction of informed traders are desirable,
and a larger fraction of order imbalances should be traded by market makers.
In all cases time priority is applied. Three multi-price rules are found desir-
able, and the selection between them is correlated with the tick size. Other
trading factors, such as the price adjustment of the market maker and the
fraction of order size revised by informed traders have ambiguous influences
on the outcomes.

B.3   Bid-ask spread and return volatility

Finally, we investigate the optimal market structure using the fitness function
of minimizing return volatility and minimizing bid-ask spread. The volatility
and inside spread of the final generation for the entire population for all 30
runs of our computer experiments restricted to the region close to the Pareto-
efficient frontier are shown in figure B.13. Based on the Pareto-efficient
frontier, which is sketched by the line to the lower right, we observe a trade-off
relationship between the two objective functions. While the return volatility
is low, the corresponding bid-ask spread is large, on the contrary, while the
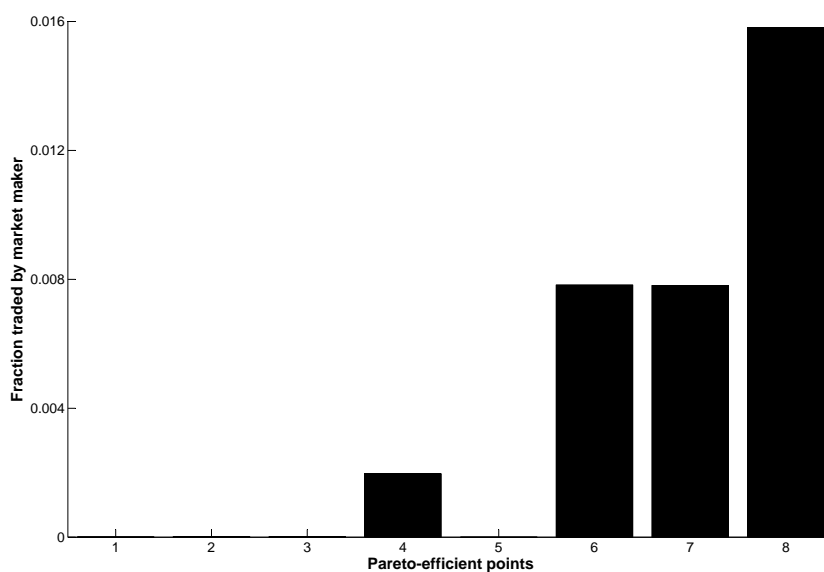
(a) the best-found fraction of order imbalance traded by market maker
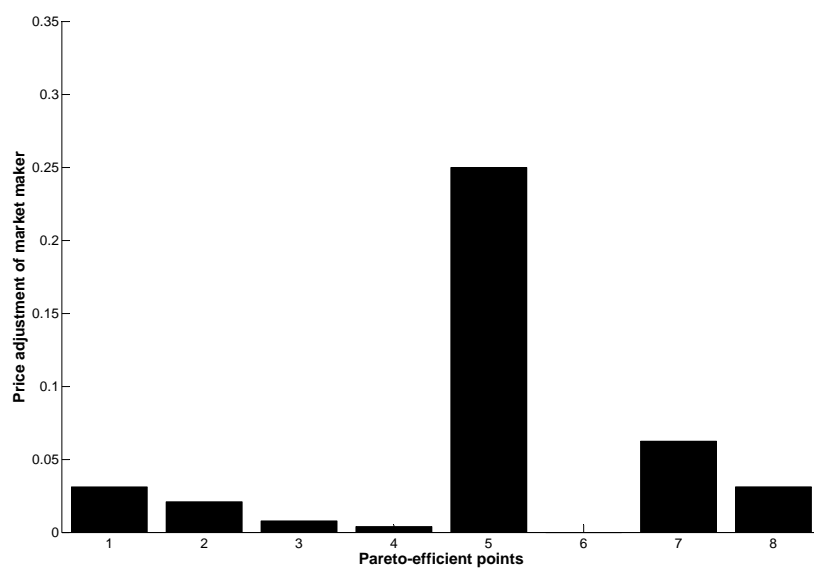
(b) the best-found price adjustment of market maker

Fig. B.12: Simulation results of intervention of market maker for maximizing trading volume and minimizing volatility

return volatility is high, the spread would be small. Due to the lack of efficient results, we only observe five markets which approximately determine the Pareto-efficient frontier, these points are identified as large points and
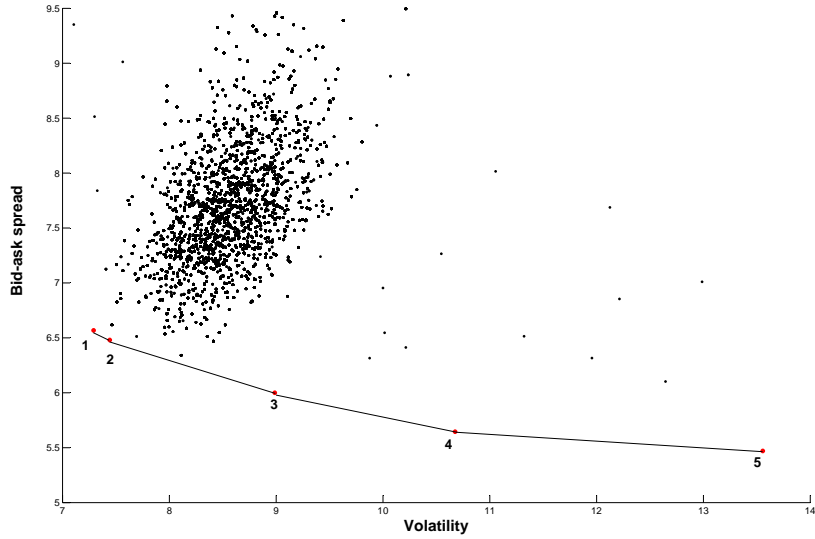
Fig. B.13: Market performance of minimizing bid-ask spread and volatility
of returns after 500 generations

associated with numbers; the market structures of these four markets are
discussed in the following section.

B.3.1   Tick size

As shown in figure B.14, for the market structures with low volatility and
large spread, the best-found tick size is relatively larger at about 3.3, and it
declines dramatically for the market structure with high volatility and narrow
spread identified by point 4. Surprisingly, this result is not consistent with
the findings from single-objective optimization experiments, in which a large
tick size is required to minimize the bid-ask spread, and a comparatively
small tick is desirable to minimize the return volatility.

B.3.2   Priority rules and multi-price rules

For the three best-found market structures having relatively low volatility
and large spread, which are represented by points 1 to 3, in case of multiple
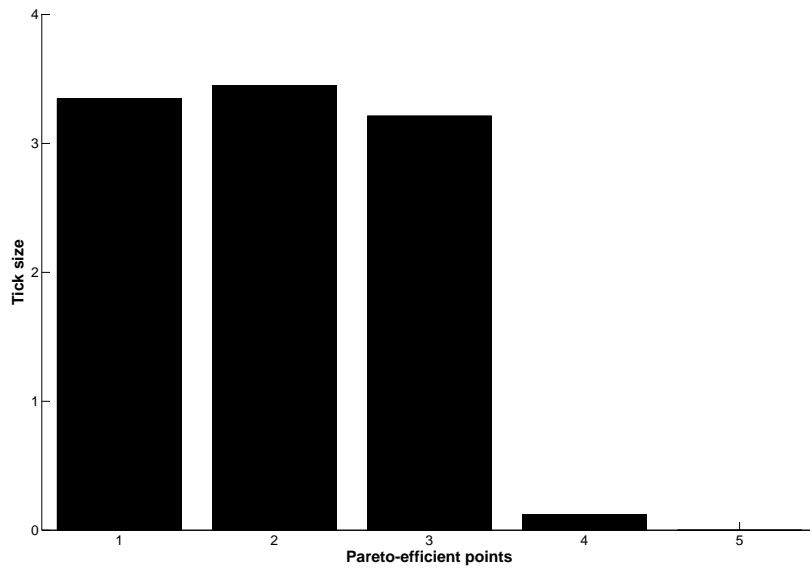
Fig. B.14: the best-found tick sizes for minimizing spread and volatility
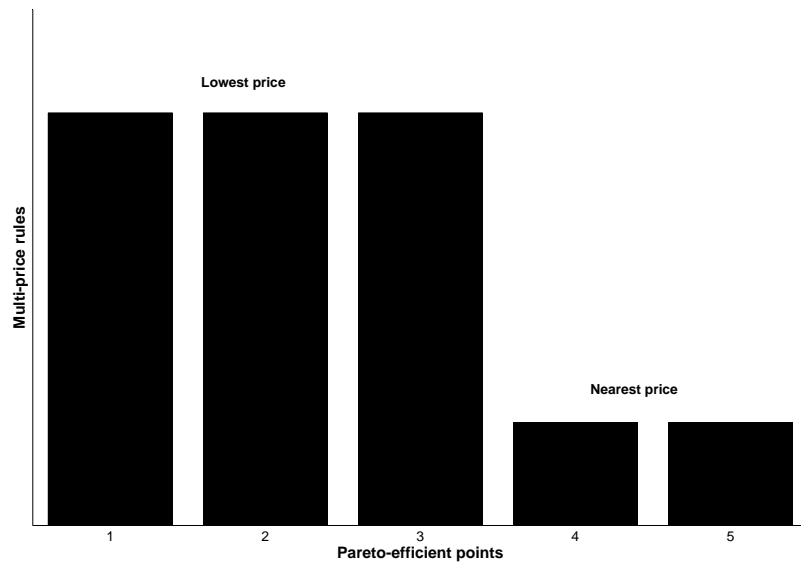
Fig. B.15: the best-found multi-price rules for minimizing spread and volatility

prices, the rules of selecting the lowest price should be applied, and the random selection priority rule is most preferred, see figure B.15 and B.16. For the market structure with small spread and high volatility, the multi-

price rule of choosing the nearest price to the previous price is obtained, and the time priority rules is used to determine the rationing of orders.
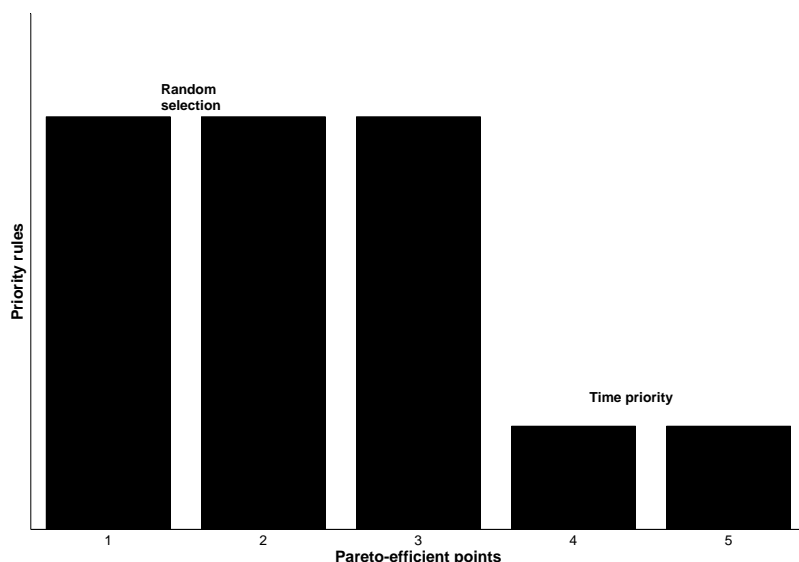
Fig. B.16: the best-found priority rules for minimizing spread and volatility

B.3.3   Transparency

Although ambiguous results of the influence of the number of informed traders on the best-found market structure have been found in the single-objective optimization experiments with objective of either minimizing spread or minimizing volatility, in this test, we find that unless the decision maker puts more emphasis on a small volatility at the expense of the spread (where only about 10% of traders are informed), the fraction of informed traders should be comparatively large that about half of traders should be informed and react immediately on any order imbalance. However, the fraction of the order size revised by those informed traders shows a decreasing trend, see figure B.3.3.
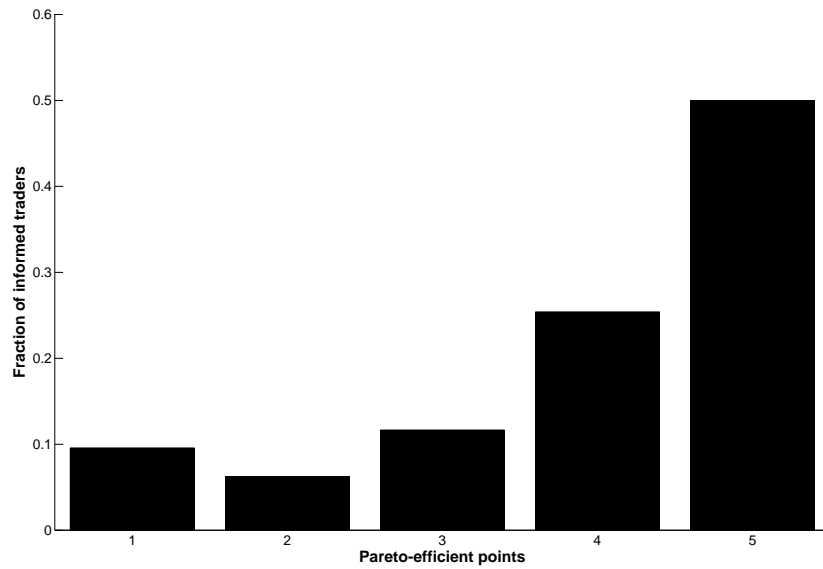
(a) the best-found fraction of informed traders for minimizing spread and volatility
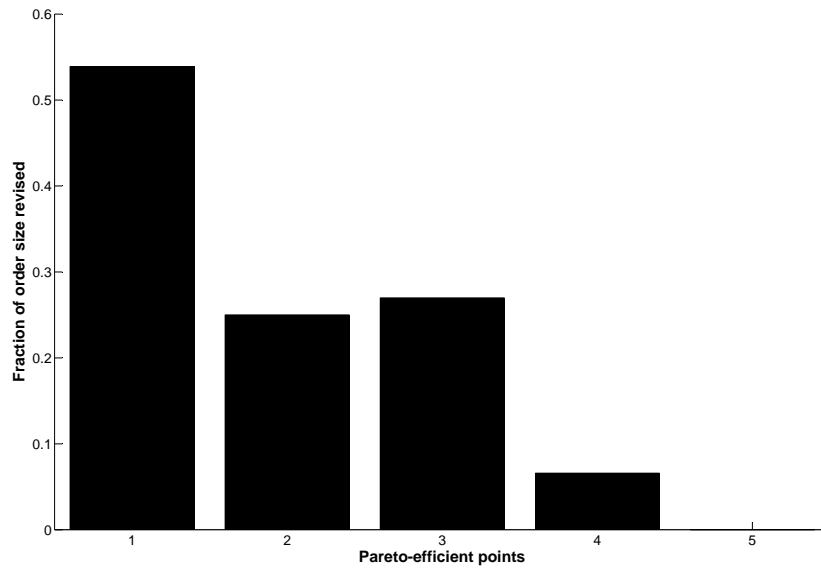
(b) the best-found fraction of order size revised for minimizing spread and volatility

Fig. B.17: Simulation results of market transparency for minimizing spread and volatility

B.3.4   The intervention of market maker

Results, shown in figure B.3.4, disclose that when concerns about the bid-ask spread are dominating a relatively higher fraction of order imbalance (about 0.39%) compared with the other best-found market structures should be traded by market maker and a smaller intervention of market maker if concerns about return volatility are more important. However, for either case, the best-found fraction is very small with an average value of approximately 0.1%. This result is consistent with the one from the previous experiments applying only a single objective function - the return volatility and bid-ask spread, respectively. In addition, we notice that while the best-found fraction traded by market maker is the highest, the price adjustment factor is maximal. It decreases significantly with less amount traded by market maker.

B.3.5   Summary of results

The main results from this computer experiment can be summarized that for market with small spread (and associated high volatility), relatively small fraction of order imbalance should be traded by market maker; larger fraction of traders should get pre-trade information; the multi-price rule of selecting price closest to the previous price and the time priority rule are applied, and a smaller tick size is required.
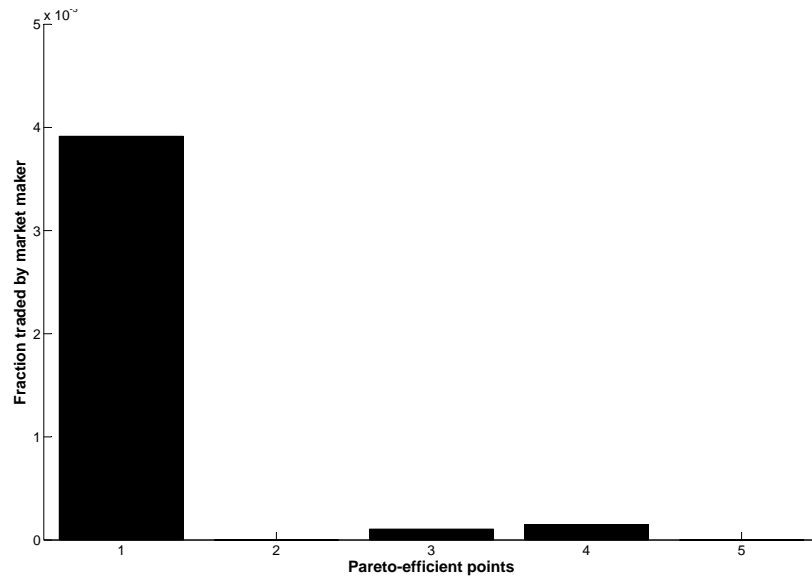
(a) the best-found fraction of order imbalance traded by market maker
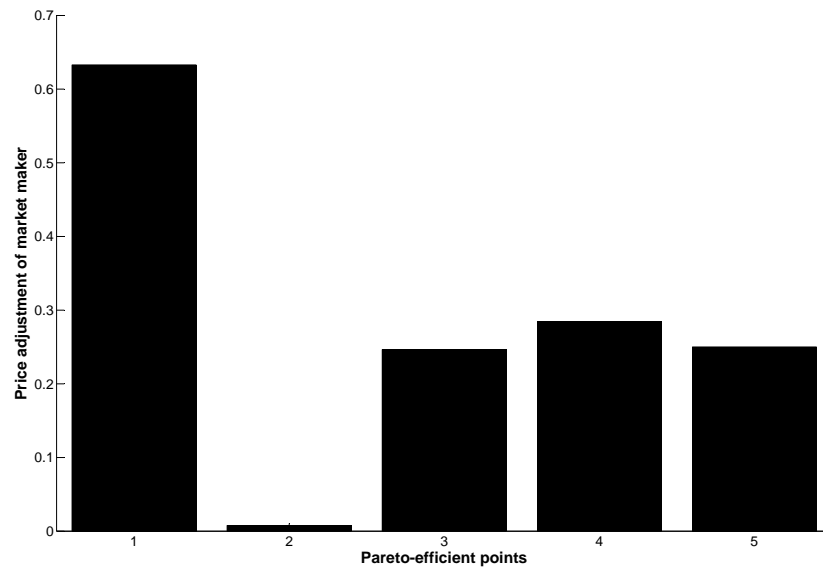
(b) the best-found price adjustment of market maker

Fig. B.18: Simulation results of intervention of market maker for minimizing spread and volatility

Bibliography

Abido, M. and J. Bakhashwain (2003). A novel multiobjective evolutionary algorithm for optimal reactive power dispatch problem. In *Proceedings of the 2003 10th IEEE International Conference on Electronics, Circuits and Systems, 2003. ICECS 2003.*, Volume 3, pp. 1054 – 1057.

Abraham, A. and R. Goldberg (Eds.) (2005). *Evolutionary Multiobjective Optimization: Theoretical Advances and Applications*. Springer.

Ackert, L. and R. Deaves (2010). *Behavioral Finance: Psychology, Decision-Making, and Markets*. South-Western Cengage Learning.

Ahn, H.-J., C. Q. Cao, and H. Choe (1996). Tick size, spread and volume. *Journal of Financial Intermediation 5*, 2–22.

Aitken, M. J., H. Berkman, and D. Mak (2001). The use of undisclosed limit orders on the australian stock exchange. *Journal of Banking & Finance 25*(8), 1589–1603.

Aittokoski, T. and K. Miettinen (2008). Efficient evolutionary method to approximate the pareto optimal set in multiobjective optimization. In *EngOpt 2008 - International Conference on Engineering Optimization*.

Alexandrova-Kabadjova, B., A. Krause, and E. P. K. Tsang (2006). Competition among payment networks using generalized population based incremental learning. In *12th International Conference on Computing in Economics and Finance (CEF2006), Limassol, Cyprus*.

Alexandrova-Kabadjova, B., E. Tsang, and A. Krause (2011). Profit-maximising strategies for an artificial payment card market, is learning

possible? *Journal of Intelligent Learning Systems and Applications 3*, 77–81.

Allaudeen, H. and T. Eric (1998). The effect of tick size on price clustering and trading volume. *Journal of Business Finance & Accounting 25*(7-8), 849–867.

Altman, M. (2006). *Handbook of Contemporary Behavioral Economics: Foundations and Development.* M.E.Sharp, Inc.

Amihud, Y. and H. Mendelson (1980). Dealership market : Market-making with inventory. *Journal of Financial Economics 8*, 31–53.

Amihud, Y. and H. Mendelson (1982). Asset pricing behavior in a dealership market. *Financial Analysts Journal 38*, 50–59.

Amihud, Y., H. Mendelson, and M. Murgia (1990). Stock market microstructure and return volatility, evidence from italy. *Journal of Banking and Finance 14*, 423–440.

Angel, J. (1997). Tick size, share prices and stock splits. *Journal of Finance 52*, 655–681.

Angel, J. and D. Weaver (1998). Priority rules! Working paper, Rutgers University.

Anirban, C., T. I. Muni, P. Marco, and A. Frederic (2011). Econophysics: Empirical facts and agent-based models. *Quantitative Finance 11*, 991–1012.

Ariely, D., A. Ockenfels, and A. Roth (2005). An experimental analysis of ending rules in internet auctions. *The RAND Journal of Economics 36*, 790–809.

Arthur, W., J. H. Holland, B. LeBaron, R. Palmer, and P. Taylor (1997). Asset pricing under endogenous expectations in an artificial stock market. In *The Economy as an Evolving Complex System II, Reading, MA: AddisonWesley,*, pp. 15–44.

Axelrod, R. (1987). The evolution of strategies in the iterated prisoner's dilemma. In *Genetic algorithms and simulated annealing, Research notes in AI*. Pitman/Morgan Kaufmann.

Ayala, H. V. H. and L. dos Santos Coelho (2008). A multiobjective genetic algorithm applied to multivariable control optimization. *ABCM Symposium Series in Mechatronics 3*, 736–745.

Bacidore, J. (1997). The impact of decimalization on market quality: An empirical investigation of the toronto stock exchange. *Journal of Financial Intermediation 6*, 92–120.

Back, K., C. H. Cao, and G. A. Willard (2000). Imperfect competition among informed traders. *Journal of Finance 55*, 2117–2155.

Back, T., U. Hammel, and H.-P. Schwefel (1997). Evolutionary computation: Comments on the history and current state. *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION 1*, 3–17.

Bagehot, W. (1971). The only game in town. *Financial Analysts Journal 22*, 12–14.

Bajari, P. and A. Hortacsu (2003). The winners curse, reserve prices, and endogenous entry: empirical insights from ebay auctions. *RAND Journal of Economics 34*, 329–355.

Bak, P., M. Paczuski, and M. Shubik (1997). Price variations in a stock market with many agents. *Physica A 246*, 430–453.

Ball, C. A. and T. Chordia (2001). True spreads and equilibrium prices. *Journal of Finance 56*, 1801–1835.

Baluja, S. (1994). Population-based incremental learning: A method for integrating genetic search based function optimization and competitive learning. Technical report cmu-cs-94-163, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA.

Baluja, S. (1995). An empirical comparison of seven iterative and evolutionary function optimization heuristics. Technical reports, Carnegie Mellon University. Available via. anonymous ftp at: reports.adm.cs.cmu.edu.CMU-CS-95-193.

Baluja, S. and R. Caruana (1995). Removing the genetics from the standard genetic algorithm. In *In: Proceedings of international conference on machine learning, Morgan Kaufmann*, pp. 38–46.

Baluja, S. and S. Davies (1997). Using optimal dependency trees for combinatorial optimization: Learning the structure of search space. Technical Report CMU-CS-97-107, Carnegie Mellon University, Pittsburgh, Pennsylvania.

Bapna, R., A. G. Paulo Goes, and Y. Jin (2004). User heterogeneity and its impact on electronic auction market design: An empirical exploration. *MIS Quarterly 28*(1), 21–43.

Barber, B. and T. Odean (2001). Boys will be boys: gender, overconfidence, and common stock investment. *Quarterly Journal of Economics 141*, 261–292.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of Financial Economics 49*, 307–343.

Barberis, N. and R. Thaler (2002). A survey of behavioral finance. Working paper, National Bureau of Economic Research.

Barberis, N. and R. Thaler (2003). *Handbook of the Economics of Finance*, Chapter A survey of behavioral finance, pp. 1053–1128. Elsevier North Holland.

Barclay, M., T. Hendershott, and D. T. McCormick (2003). Competition among trading venues: information and trading on electronic communications networks. *Journal of Finance 58*, 2637–2665.

Barth III, M. J., E. M. Remolona, and P. D. Wooldridge (2002). Changes in market functioning and central bank policy: An overview of the issues. In *Bank for International Settlements Working Paper No. 120*.

Baruch, S. (1997). Who benefits from an open limit order book? Working Paper, Washington University.

Becker, G. S. (1962). Irrational behavior and economic theory. *Journal of Political Economy 70*, 1–13.

Benhamou, E. and T. Serval (2000). On the competition between ecns, stock markets and market makers. In *Electronic Commerce and Web Technologies*, Volume 1875 of *Lecture Notes in Computer Science*, pp. 291–300. Springer Berlin Heidelberg.

Bererton, C., G. Gordon, and S. Thrun (2003). Auction mechanism design for multi-robot coordination. In *Proceedings of 17th Annual Conf. on Neural Information Processing Systems (NIPS03)*.

Bessembinder, H. (2000). Tick size, spreads, and liquidity: An analysis of nasdaq securities trading near ten dollars. *Journal of Financial Intermediation 9*(3), 213–239.

Bessembinder, H., W. Maxwell, and K. Venkataraman (2006). Market transparency, liquidity externalities, and institutional trading costs in corporate bonds. *Journal of Financial Economics 82*, 251–288.

Biais, B. (1993). Price formation and equilibrium liquidity in fragmented and centralized markets. *Journal of Finance 48*, 157–184.

Biais, B., L. Glosten, and C. Spatt (2005). Market microstructure: A survey of microfundations, empirical results, and policy implications. *Journal of Financial Markets*, 217–264.

Biais, B., P. Hillion, and C. Spatt (1995). An empirical analysis of the limit order book and the order flow in the paris bourse. *Journal of Finance 50*, 1655–1689.

Biais, B., P. Hillion, and C. Spatt (1999). Price discovery and learning during the preopening period in the paris bourse. *Journal of Political Economy 107*, 1218–1247.

Bin, Z., Y. Maosong, and X. Xianya (2004). The comparisons between pricing methods on pool-based electricity market using agent-based simulation. In *Proceedings of the IEEE International Conference on Electric Utility Deregulation, Restructuring and Power Technologies (DRPT2004),*, Hong Kong.

Binmore, K. and P. Klemperer (2002). The biggest auction ever: the sale of the british 3g telecom licences. *Economic Journal 112*, C74–C96.

Bloomfield, R. and M. OHara (1999). Market transparency: who wins and who loses? *Review of Financial Studies 12*, 5–35.

Bloomfield, R. and M. OHara (2000). Can transparent markets survive? *Journal of Financial Economics 55*, 259–425.

Board, J. and C. Sutcliffe (1995). The e!ects of trade transparency in the london stock exchange: a summary. Working Paper, London School of Economics.

Boehmer, E., G. Saar, and L. Yu. (2005). Lifting the veil: An analysis of pre-trade transparency at the nyse. *Journal of Finance 60*, 783–815.

Bondarenko, O. (2001). Competing market makers,liquidity provision,and bidask spreads. *Journal of Financial Markets 4*, 269–308.

Bondt, W. F. M. D. and R. Thaler (1985). Does the stock market overreact. *The Journal of Finance 40*, 793–805.

Borgers, T. and C. Dustmann (2002). Rationalizing the umts spectrum bids: The case of the uk auction. *ifo Studien 48*(1).

Bower, J. and D. Bunn (2000). Model-based comparison of pool and bilateral markets for electricity. *Energy Journal 21*(3), 1–29.

Bower, J., D. Bunn, and C. Wattendrup (2001). A model-based analysis of strategic consolidation in the german electricity industry. *Energy Policy 29*, 987–1005.

Brabazon, A. and M. O'Neill (2003a, 8 September, 2003). Evolving financial models using grammatical evolution. In *In Proceedings of The Annual Conference of the South Eastern Accounting Group (SEAG)*, London Metropolitan University, London.

Brabazon, A. and M. O'Neill (2003b, June 2003). A grammar model for foreign exchange trading. In *Proceedings of the International Conference on Artificial Intelligence*, pp. 492498.

Branke, J., T. Kaussler, C. Schmidt, and H. Schmeck (2000). A multi-population approach to dynamic optimization problems. In *In Proceedings of the 4th International Conference on Adaptive Computing in Design and Manufacturing*.

Breidt, F. J., N. Crato, and P. D. Lima (1998). The detection and estimation of long memory in stochastic volatility. *Journal of Econometrics 83*, 325–348.

Bruce, B. (2010). *Handbook of Behavioral Finance*. Edward Elgar Publishing Limited.

Brusco, S. and M. Jackson (1999). The optimal design of a market. *Journal of Economic Theory 88*, 1–39.

Bunn, D. and C. Day (1998, 1821 October). Agent-based simulation of electric power pools: a comparison with the supply function equilibrium approach. In *in: Technologys Critical Role in Energy and Environmental Markets, Proceedings of the 19th Annual North American Conference of the United States Association for Energy Economics and the International Association for Energy Economics,*, Albuquerque, New Mexico. IAEE/USAEE, Cleveland.

Bunn, D. and F. Oliveira (2003). Evaluating individual market power in electricity markets via agent-based simulation. *Annals of Operations Research 121*(1-4), 57–77.

Bureerat, S. and S. Srisomporn (2010). Optimum plate-fin heat sinks by using a multi-objective evolutionary algorithm. *Engineering Optimization 42*, 305–323.

Bureerat, S. and K. Sriworamas (2007). Population-based incremental learning for multiobjective optimisation. In *Soft Computing in Industrial Applications*, Volume 39 of *Advances in Soft Computing*, pp. 223–232. Springer Berlin / Heidelberg.

Byde, A. (2002). Applying evolutionary game theory to auction mechanism design. Technical report, Hewlett-Packard Technical Report HPL-2002-321.

Calamia, A. (1999). Market microstructure: Theory and empirics. Lem papers series, Laboratory of Economics and Management (LEM), Sant'Anna School of Advanced Studies, Pisa, Italy.

Camerer, C. F., G. Loewenstein, and M. Rabin (2004). *Advances in Behavioral Economics*. Princeton University Press.

Campbell, M. I., J. Cagan, and K. Kotovsky (1999). A-design: An agent-based approach to conceptual design in a dynamic environment. *Research in Engineering Design 11*(3), 172–192.

Carceles-Poveda, E. and C. Giannitsarou (2008). Asset pricing with adaptive learning. *Review of Economic Dynamics 11*, 629–651.

Cau, T. (2003). *Analyising Tacit Collusion in Oligopolistic Electricity Markets Using a Co-evolutionary Approach*. Ph. D. thesis, Australian Graduate School of Management.

Cau, T. and E. Anderson (2002). A co-evolutionary approach to modeling the behavior of participants in competitive electricity markets. In *in: Pro-*

ceedings of the Power Engineering Society Summer Meeting, pp. 15341540. IEEE Society Press, Piscataway, N.J.

Chakravarty, S. and C. W. Holden (1995). An integrated model of market and limit orders. *Journal of Financial Intermediation 4*, 213–241.

Challet, D. and R. Stinchcombe (2001). Analyzing and modeling 1+1d markets. *Physica A 300*, 285–299.

Chan, K. C. and C. yang Hwang (1998). The impact of tick size on market quality: an empirical investigation of the stock exchange of hong kong. working paper, Hong Kong University of Science & Technology.

Chen, Shu-Heng, e. (Ed.) (2002). *Evolutionary Computation in Economics and Finance, volume 100 of Studies in Fuzziness and Soft Computing.* Physica-Verlag Heidelberg. ISBN: 3-7908-1476-8.

Chen, S.-H. and B.-T. Chie (2008). Lottery markets design, micro-structure, and macro-behavior: An ace approach. *Journal of Economic Behavior & Organization 67*(2), 463 – 480.

Chiarella, C. and G. Iori (2002). A simulation analysis of the microstructure of double auction markets. *Quantitative Finance 2*, 346–353.

Choi, J., G. G. Lim, and K. C. Lee (2010). An agent based market design methodology for combinatorial auctions. *Journal of Artificial Societies and Social Simulation 13*, 2.

Chowdhry, B. and V. Nanda (1991). Multi-market trading and market liquidity. *Review of Financial Studies 4*, 483–511.

Chun, S. H. and R. J. La (2009). Auction mechanism for spectrum allocation and profit sharing. In *Proc. Int. Conf. Game Theory for Networks GameNets '09*, pp. 498–507.

Cincotti, S. and S. Pastore (2005). Heterogeneous information-based artificial stock market. In *WEHIA (13-15 June 2005, Essex, United Kingdom).*

Cincotti, S., L. Ponta, and M. Raberto (2005). A multi-assets artificial stock market with zero-intelligence traders. In *WEHIA (13-15 June 2005, Essex, United Kingdom)*.

Cliff, D. (2001, May 1998). Evolutionary optimization of parameter sets for adaptive software-agent traders in continuous double-auction markets. In *in: Presented at the Artificial Societies and Computational Markets (AS-CMA98) workshop at the Second International Conference on Autonomous Agents*. Also available as HP Labs Technical Report HPL-2001-99.

Cliff, D. (2002a). Evolution of market mechanism through a continuous space of auction-types ii: Two-sided auction mechanisms evolve in response to market shocks. In *Proceedings of the International Conference on Internet Computing IC02, Vol III, CSREA Press*, Washington, DC, pp. 682688. Also available as HewlettPackard Laboratories Technical Report HPL-2002-128.

Cliff, D. (2002b, December 2002). Visualizing search-spaces for evolved hybrid auction mechanisms. In *Presented at the Beyond Fitness: Visualizing Evolution workshop at the 8th International Conference on the Simulation and Synthesis of Living Systems (ALifeVIII) conference*, Sydney. Also available as HewlettPackard Laboratories Technical Report HPL-2002-291.

Cliff, D. (2003). Explorations in evolutionary design of online auction market mechanisms. *Electronic Commerce Research and Applications 2*, 162–175.

Cliff, D. and J. Bruten (1997). Minimal intelligence agents for bargaining behaviors in market-based environments. Technical Report HPL-97-91, Hewlett Packard Laboratories, Bristol, England.

Coello Coello, C. A. (1996). *An empirical study of evolutionary techniques for multiobjective optimization in enginerring design*. Ph. D. thesis, Department of Computer Science, Tulance University.

Coello Coello, C. A. (1998). A comprehensive survey of evolutionary-based multiobjective optimization techniques. *Knowledge and Information Systems 1*, 269–308.

Coello Coello, C. A., G. B. Lamont, and D. A. V. Veldhuizen (2002). *Evolutionary algorithms for solving multi-objective problems.* Kluwer Academic Publishers, Bolston.

Cohen, K., R. Conroy, and S. Maier (1985). Order flow and the quality of the market. In Y. Amihud, T. Ho, and R. A. Schwartz (Eds.), *Market Making and the Changing Structure of the Securities Industry*, pp. 93–110. Rowman & Littlefield, Lanham.

Cohen, K., S. Maier, R. Schwartz, and D. Whinston (1981). Transaction costs, order placement strategy, and existence of the bid-ask spread. *Journal of Political Economy 89*, 287–305.

Cohen, K., S. Maier, R. A. Schwartz, and D. K. Whitcomb (1978). Limit orders, market structure, and the returns generation process. *Journal of Finance 33*, 723–736.

Comerton-Forde, C., S. Lau, and T. McInish (2005). Opening and closing behavior following the introduction of call auctions in singapore. Working Paper.

Comerton-Forde, C. and K. Tang (2009). Anonymity, liquidity and fragmentation. *Journal of Financial Markets 12*, 337–367.

Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance 1*, 223–236.

Cont, R. (2007). Volatility clustering in financial markets: Empirical facts and agent-based models. In G. Teyssiere and A. P. Kirman (Eds.), *Long Memory in Economics*, Chapter 10, pp. 289–309. Springer Berlin Heidelberg.

Cont, R., M. Potters, and J.-P. Bouchaud (1997). Scaling in financial data: Stable laws and beyond. In *Scale Invariance and Beyond*, Chapter Scaling in financial data: Stable laws and beyond. Springer, Berlin.

Copeland, T. and D. Galai (1983). Information effects on the bid-ask spread. *Journal of Financial Economics 8*, 31–53.

Cordella, T. and T. Foucault (1996). Minimum price variations, time priority, and quote dynamics. Working paper, Universitat Pompeu Fabra.

Cordella, T. and T. Foucault (1999). Minimum price variations, time priority and quote dynamics. *Journal of Financial Intermediation 8*, 141–173.

Costa, L. and P. Oliveira (2002). An evolution strategy for multiobjective optimization. *Proceedings of the World on Congress on Computational Intelligence 1*, 97–102.

Costa, L. and P. Oliveira (2004). An elitist genetic algorithm for multiobjective optimization. In *MIC 2001 - 4th Metaheuristics International Conference*, Norwell, MA, USA, pp. 217–236. Kluwer Academic Publishers.

Cramton, P. (1998). Ascending auctions. *European Economic Review 42*, 745–756.

Cramton, P., Y. Shoham, and R. Steinberg (Eds.) (2006). *Combinatorial Auctions*. MIT Press.

Cross, R., M. Grinfeld, H. Lamba, and T. Seaman (2005). A threshold model of investor psychology. *Physica A 354*, 463–478.

Daniel, K., D. Hirshleifer, and A. Subrahmanyam (1998). Investor psychology and security market under- and overreactions. *Journal of Finance 53*(6), 1839–1885.

Darley, V., A. Outkin, T. Plate, and F. Gao (2000). Sixteenths or pennies? observations from a simulation of the nasdaq stock market. Technical report, Technical report, Santa Fe Institute.

Das, R., J. Hanson, J. Kephart, and G. Tesauro (2001). Agenthuman interactions in the continuous double auction. In *in: Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI-01)*.

David, E., R. A, S. J, K. S, and J. NR (2005). Optimal design of english auctions with discrete bid levels. In *Proceedings of the 6th ACM COnference on Electronic Commerce (EC05), Vancouver, Canada*, pp. 98–107.

Davidson, E. M. and S. D. J. McArthur (2005). Concepts and approaches in multi-agent systems for power applications. In *Proceedings, Intelligent System Applications in Power (ISAP) Conference, Washington D.C.*, pp. 391–395.

Day, C. D. (2006). Application of an evolutionary algorithm to multivariate optimal allocation in stratified sample designs. In *American Statistical Association Conference, Joint Statistical Meeting 2006*.

Day, R. H. and W. H. Huang (1990). Bulls, bears, and market sheep. *Journal of Economic Behavior and Organization 14*, 299–330.

De Bonet, J. S., C. L. Isbell, and P. Viola (1997). Mimic: Finding optima by estimating probability densities. In *Proceedings of NIPS'97*, pp. 424431. MIT Press, Cambridge, MA.

de Vries, S. and R. V. Vohra (2003). Combinatorial auctions: A survey. *Informs Journal on Computing 15*, 284–309.

Deb, K. (1998). Multi-objective genetic algorithms: problem difficulties and construction of test problems. Technical report ci-49/98, Department of Computer Science, University of Dortmund, Germany.

Deb, K. (1999, September). Multi-objective genetic algorithms: Problem difficulties and construction of test problems. *Evolutionary Computation 7*, 205–230.

Deb, K. (2001). *Multi-Objective Optimization using Evolutionary Algorithm*. John Wiley & Sons, Ltd.

Deb, K. (2008). Introduction to evolutionary multiobjective optimization. In *Multiobjective Optimization*, pp. 59–96.

Deb, K., A. Pratap, S. Agarwal, and T. Meyarivan (2002). A fast and elitist multi-objective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation 6*, 182–197.

Demarchi, M. and T. Foucault (2000). Equity trading systems in europe, a survey of recent changes. *Journal of Economics and Statistics*, 60.

Demsetz, H. (1968). The cost of transacting. *Quarterly Journal of Economics 82*, 33–53.

Domowitz, I. (1993). A taxonomy of automated trade execution systems. *Journal of International Money and Finance 12*, 607–631.

Duffy, J. and M. Unver (2008). Internet auctions with artificial adaptive agents: A study on market design. *Journal of Economic Behavior & Organization 67*(2), 394 – 417.

Easly, D. and M. OHara (1987). Price, trade size, and information in securities markets. *Journal of Financial Economics 19*, 69–90.

Economides, N., R. S. (1995). Electronic call market trading. *Journal of Portfolio Management 21*, 10–18.

Edgeworth, F. (1881). *Mathematical Physics*. P. Keagan, London, England.

Edwards, A., L. Harris, and M. Piwowar (2007). Corporate bond market transaction costs and transparency. *Journal of Finance 62*, 1421–1451.

Ehrentreich, N. (2007). *Agent-Based Modeling: The Santa Fe Institute Artificial Stock Market Model Revisited*. Springer.

Eiben, A. E., J. E. Smith, and F. Cover (2003). *Introduction to evolutionary computing*. Springer.

Ellul, A., C. Holden, P. Jain, and R. Jennings (2002). Determinates of order choice on the new york stock exchange. Working paper, Indiana University.

Erev, I. and A. Roth (1998). Predicting how people play games: reinforcement learning in experimental games with unique mixed strategy equilibria. *American Economic Review 88*, 848–881.

Erickson, M., A. Mayer, and J. Horn (2001). The niched pareto genetic algorithm 2 applied to the design of groundwater remediation systems. In *Evolutionary Multi-Criterion Optimization*, Volume 1993 of *Lecture Notes in Computer Science*, pp. 681–695. Springer Berlin / Heidelberg.

Ewerhart, C. and B. Moldovanu (2002). The german umts design: Insights from multi-object auction theory. CESifo Working Paper Series 680, CESifo Group Munich.

Fernandez-Blanco, P., D. J. Bodas-Sagi, F. J. Soltero, and J. I. Hidalgo (2008). Technical market indicators optimization using evolutionary algorithms. In *Proceedings of the 2008 GECCO conference companion on Genetic and evolutionary computation*.

Fisher, I. (1930). *The Theory of Interest*. (MacMillan, New York.

Fishman, M. and F. A. Longstaff (1992). Dual trading in futures markets. *Journal of Finance 47*, 643–671.

Flood, M. D., R. Huisman, K. Koedijk, and R. J. Mahieu (1999). Quote disclosure and price discovery in multiple-dealer finance markets. *Review of Financial Studies 12*, 37–59.

Flood, Mark D., R. H. K. G. K. and R. J. Mahieu (1999). Quote disclosure and price discovery in multiple dealer financial markets. *Review of Financial Studies 12*, 37–59.

Fogel, D. (1995). *Evolutionary Computation: Toward a New Philosophy of Machine Intelligence*. Piscataway, NJ: IEEE Press.

Fogel, D. B. (1994). An introduction to simulated evolutionary optimization. *IEEE Trans. on Neural Networks: Special Issue on Evolutionary Computation 5*, 3–14.

Fonseca, C. M. and P. J. Fleming (1993). Genetic algorithms for multiobjective optimization: Formulation, discussion, and generalization. In *In: Proceedings of the Fifth International Conference on Genetic Algorithms*, pp. 416–423.

Fonseca, C. M. and P. J. Fleming (1995a). Multiobjective genetic algorithms made easy: Selection, sharing and mating restrictions. In *In First International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications (GALESIA 95)*, pp. 45–52.

Fonseca, C. M. and P. J. Fleming (1995b, March). An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation 3*, 1–16.

Fonseca, C. M. and P. J. Fleming (1998). Multiobjective optimization and multiple constraint handling with evolutionary algorithms-part i: A unified formulation. *IEEE Transactions on Systems, Man, and Cybernetics, Part A: Systems and Humans 28*, 26–37.

Forster, M. and T. George (1992). Anonymity in securities markets. *Journal of Financial Intermediation 2*, 168–206.

Foucault, T. (1999). Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial Markets 2*, 99–134.

Foucault, T., O. Kadan, and E. Kandel (2005). Limit order book as a market for liquidity. *Review of Financial Studies*, 1171–1217.

Foucault, T., S. Moinas, and E. Theissen (2003). Does anonymity matter in electronic limit order markets. *Review of Financial Studies 20*, 1707–1748.

Frankfurter, G. M. and E. G. McGoun (2010). Market efficiency or behavioral finance: The nature of the debate. *Journal of Psychology and Financial Markets 1*, 200–210.

Frino, A., G. Dionigi, and L. Andrew (2008). Limit order book, anonymity and market liquidity: evidence from the sydney futures exchange. *Accounting & Finance 48*, 561–573.

Frino, A., P. Forrest, and M. Duffy (1999). Life in the pits: competitive market making and inventory controlfurther australian evidence. *Journal of Multinational Financial Management 9*, 373–385.

Gandhi, S., C. Buragohain, L. Cao, H. Zheng, and S. Suri (2007). A general framework for wireless spectrum auctions. In *Proc. 2nd IEEE Int. Symp. New Frontiers in Dynamic Spectrum Access Networks DySPAN 2007*, pp. 22–33.

Garman, M. (1976). Market microstructure. *Journal of Financial Economics 3*, 257–275.

Gaunersdorfer, A. and C. H. Hommes (2001). Nonlinear adaptive beliefs and the dynamics of financial markets: The role of the evolutionary fitness measure. In *Artificial neural networks–ICANN 2001: International Conference, Vienna, Austria, August 21-25, 2001 : proceedings*.

Gemmill, G. (1994). Transparency and liquidity: a study of block trades on the london stock exchange under di!erent publication rules. *Journal of Finance 51*, 1765–1790.

Gemmill, G. (1996). Transparency and liquidity: a study of block trades on the london stock exchange under different publiation rules. *Journal of Finance 51*, 1765–1786.

Glosten, L. R. and P. Milgrom (1985). Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics 14*, 71–100.

Gode, D. K. and S. Sunder (1993). Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy 101*, 119–137.

Goel, T. and N. Stander (2009). A study on the convergence of multiobjective evolutionary algorithms. In *13th Multidisciplinary Analysis and Optimization Conference (MAO-2009)*.

Goldberg, D. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley.

Goldberg, J. and R. von Nitzsch (2001). *Behavioral Finance.* John Wiley and Sons Ltd.

Goldstein, M. A., E. S. Hotchkiss, and E. R. Sirri (2007). Transparency and liquidity: A controlled experiment on corporate bonds. *The review of financial studies 20*(2), 235–273.

Goldstein, M. A. and K. A. Kavajecz (2000). Eighths, sixteenths, and market depth: Changes in tick size and liquidity provision on the nyse. *Journal of Financial Economics 56*, 125–149.

Gosling, T., N. Jin, and E. Tsang (2004). Population based incremental learning versus genetic algorithms: Iterated prisoners dilemma. Technical report, Technical Report, University of Essex, England.

Grammig, J., D. Schiereck, and E. Theissen (2001). Knowing me, knowing you: Trader anonymity and informed trading in parallel markets. *Journal of Financial Markets 4*, 385–412.

Grinblatt, M., S. Titman, and R. Wermers (1995). Momentum investment strategies, portfolio performance, and herding: A study of mutual fund behavior. *The American Economic Review 85*, 1088–1105.

Gu, M. (1995). Market mediating behavior: An economic analysis of the security exchange specialists. *Journal of Economic Behavior and Organization 27*, 237–256.

Hailu, A. and S. Schilizzi (2004). Are auctions more efficient than fixed price schemes when bidders learn. *Australian Journal of Management 29*, 147–168.

Hailu, A. and S. Thoyer (2006). Multi-unit auction format design. *Multi-unit auction format design 1*(2), 129–146.

Hakansson, N. H., A. Beja, and J. Kale (1985). On the feasibility of auto-mated market making by a programmed specialist. *Journal of Finance 40*, 1–20.

Hamalainen, R., E. Kettunen, and H. Ehtamo (1997, Feb. 17-19). Game modeling and coordination processes for two-level multi-agent systems. In *Proc. of the 16th IASTED International Conference: Modelling, Identification, and Control*, Innsbruck.

Hamalainen, R. and J. Parantainen (1995, December 6-8). Load analysis by agent-based simulation of the electricity distribution system. In *Proc. of the 2nd IFAC symposium on control of power plants and power systems*, Number 1, Cancun, Mexico, pp. 213–217. SIPOWER95.

Hanne, T. (1999). On the convergence of multiobjective evolutionary algorithms. *European Journal of Operational Research 117*, 553–564.

Harik, G. (1999). Linkage learning via probabilistic modeling in the ecga. Technical report, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL, IlliGAL Report No. 99010, 1999.

Harik, G. R., F. G. Lobo, and D. E. Goldberg (1999). The compact genetic algorithm. *IEEE Transactions on Evolutionary Computation 3*(4), 287–297.

Harris, L. (1990a, September). Estimation of stock price variances and serial covariances from discrete observations. *Journal of Financial and Quantitative Analysis 25*(03), 291–306.

Harris, L. (1990b). Liquidity, trading rules and electronic trading systems. *New York University Salomon Centre Monograph Series in Finance*.

Harris, L. (1991). Stock price clustering and discreteness. *Review of Financial Studies 4*, 389–416.

Harris, L. (1996). Does a large minimum price variation encourage order exposure? Working paper, University of Southern California.

Harris, L. (1997). Order exposure and parasitic traders. Paper presented at the Deusche Bourse AG Symposium Equity Market Structure for Large- and Mid Cap Stocks, Frankfurt.

Harris, L. (1998). Optimal dynamic order submission strategies in some stylized trading problems. *Financial Markets, Institutions & Instruments*.

Harris, L. and M. Piwowar (2006). Secondary trading costs in the municipal bond market. *Journal of Finance 61*, 1361–1397.

Harris, L. E. (1994). Minimum price variations, discrete bid-ask spreads, and quotation sizes. *Review of Financial Studies 7*(1), 149–178.

Harris, L. E. and J. Hasbrouck (1996). Market vs. limit orders: the super dot evidence on order submission strategy. *Journal of Financial and Quantitative Analysis 31*, 213–231.

Hasker, K., R. Gonzales, and R. C. Sickles (2004). An analysis of strategic behavior and consumer surplus in ebay auctions. Working Paper, Rice University.

Heidle, H. and R. Huang (2002). Information-based tradingin dealer and auction markets: ananalysis of exchange listings. *Journal of Financial and Quantitative Analysis 37*, 391424.

Hendershott, T. (2003). Electronic trading in financial markets. *IT Professional 5*, 10–14.

Hendershott, T. and C. Jones (2005). Island goes dark: transparency, fragmentation, and regulation. *Review of Financial Studies 18*(3), 743–793.

Hirshleifer, D. (2001). Investor psychology and asset pricing. *Journal of Finance 56*(4), 1533–1597.

Ho, T. and H. Stoll (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics 9*, 47–73.

Ho, T. S. Y. and H. R. Stoll (1983). The dynamics of dealer markets under competition. *Journal of Finance 38*(4), 1053 – 1074.

Ho, T. Y. and H. R. Stoll (1980). On dealer markets under competition. *Journal of Finance 35*, 259–267.

Holland, J. (1975). *Adaptation in Natural and Artificial Systems*. University of Michigan Press.

Hommes, C. (2002). Modeling the stylized facts in finance through simple nonlinear adaptive systems. In *Proceedings of the National Academy of Sciences*, pp. 7221–7228.

Hommes, C. (2005). Heterogeneous agent models in economics and finance. Tinbergen Institute Discussion Paper, University of Amsterdam, and Tinbergen Institute.

Hoppe, H. C., P. Jehiel, and B. Moldovanu (2006). License auctions and market structure. *Journal of Economics & Management Strategy 15*(2), 371–396.

Horn, J. (1997). *Handbook of evolutionary computation*, Chapter Multicriterion Decision Making, pp. F1.9:1 – F1.9:15. UK: IOP Publishing Ltd.

Horn, J. and N. Nafpliotis (1993). Multiobjective optimization using the niched pareto genetic algorithm. Illigal report 93005, Illinois Genetic Algorithms Laboratory, University of Illinois, Urbana, Champaign.

Huang, R. D. and H. R. Stoll (1999). Tick size, bid-ask spreads and market structure. Owen Graduate School of Business, Vanderbilt University, Working paper.

Huang, R. D. and H. R. Stoll (2001). Tick size, bid-ask spreads, and market structure. *Journal of Financial and Quantitative Analysis 36*(4), 503–522.

Hunsberger, L. and B. J. Grosz (2000). A combinatorial auction for collaborative planning. In *Proceedings of 4th International Conference on Multi-Agent Systems (ICMAS-00) [icm00]*, pp. 151158.

Ingber, L. (1996). Trading markets with canonical momenta and adaptive simulated annealing. Technical Report 1996:TMCMASA, Lester Ingber Research, McLean, VA.

Ingber, L. (1999). A simple options training model. *Mathematical Computer Modelling 30*, 167–182.

Iori, G. (2002). A microsimulation of traders activity in the stock market: the role of heterogeneity, agents' interactions and trade frictions. *Journal of Economic Behavior & Organization 49*, 269–285.

Ivanova-Stenzel, R. and T. Salmon (2004). Bidder preferences among auction institutions. *Economic Inquiry 42*, 223–236.

Jackson, M. O. (2003). Efficiency and information aggregation in auctions with costly information. *Review of Economic Design 8*(2), 121–142.

Jehiel, P. and B. Moldovanu (2000). License auctions and market structure. CEPR Discussion Papers 2530, C.E.P.R. Discussion Papers.

Jehiel, P., B. Moldovanu, M. Ottaviani, and C. Propper (2003). An economic perspective on auction. *Economic Policy 36*, 269–308.

Jones, C. M. and M. L. Lipson (2001, February). Sixteenths: direct evidence on institutional execution costs. *Journal of Financial Economics 59*(2), 253–278.

Kanyakam, S. and S. Bureerat (2007). Passive vibration suppression of a walking tractor handlebar structure using multiobjective pbil. In *IEEE Congress on Evolutionary Computation, 2007. CEC 2007.*, pp. 4162–4169.

Kanyakam, Siwadol, S. S. and S. Bureerat (2009). Optimal geometrical design of multiple heights pin-fin heat sink using mopbil. In *The 23rd Conference of the Mechanical Engineering Network of Thailand.*

Karolyi, G. A. (2006). The world of cross-listings and cross-listings of the world: Challenging conventional wisdom. *Review of Finance (2006) 10 (1): 99-152. 10 (1)*, 99–152.

Kavejecz, K. (1999). A specialist's quoted depth and the limit order book. *Jolurnlal of Finiance 54*, 747–771.

Kazancioglu, E., G. Wu, J. Ko, S. Bohac, Z. Filipi, S. J. Hu, D. Assanis, and K. Saitou (2003). Robust optimization of an automobile valvetrain using a multiobjective genetic algorithm. In *Proceedings of DETC'03 ASME 2003 Design Engineering Technical Conferences Chicago, Illinois, September 2-6, 2003*.

Ke, M., C. Jiang, and Y. Huang (2004). The impact of tick size on intraday stock price behaviour: evidence from the taiwan stock exchange. *Pacific-Basin Finance Journal 12*, 19 39.

Keeney, R. L. and H. Raiffa (1976). *Decisions with multiple objectives: preferences and value tradeoffs.* Wiley, NY.

Keynes, J. M. (1936). *The General Theory of Employment, Interest and Money.* Macmillan Cambridge University Press.

Kim, G. and H. Markowitz (1989). Investment rules, margin, and market volatility. *Journal of Portfolio Management 16*, 45–52.

Kim, J.-H., Y.-H. Kim, S.-H. Choi, and I.-W. Park (2009). Evolutionary multi-objective optimization in robot soccer system for education. *IEEE COMPUTATIONAL INTELLIGENCE MAGAZINE 4*, 31–41.

Kim, J.-L. (2010). Examining the relationship between algorithm stopping criteria and performance using elitist genetic algorithm. In *Proceedings of the 2010 Winter Simulation Conference*.

Kittsteiner, T. and A. Ockenfels (2006). Market design: a selective review. In *Zeitschrift fur Betriebswirtschaft, Special Issue*, Volume 5, pp. 121–143.

Klemperer, P. (2002). Some observations on the british and german 3g telecom auctions. CEPR Discussion Papers 3605, C.E.P.R. Discussion Papers.

Klemperer, P. (2004). *Auctions: theory and practice.* Princeton University Press, Princeton, US.

Klemperer, P. D. and M. A. Meyer (1989). Supply function equilibria in oligopoly under uncertainty. *Econometrica: Journal of the Econometric Society 57*(6), 1243–1277.

Knowles, J. and D. Corne (2000). Approximating the non-dominated front using the pareto archived evolution strategy. *Evolutionary Computation 8*, 149–172.

Koersrindartoto, D. (2005, 1013 June). Bidding rules in deregulated day-ahead electricity markets. In *paper presented at the Workshop on the Economics of Heterogeneous Interacting Agents (WEHIA2005)*, Colchester, UK.

Koesrindartoto, D., J. Sun, and L. Tesfatsion (2005). An agent-based computational laboratory for testing the economic reliability of wholesale power market desings. In *Proceedings of the IEEE Power Engineering Society General Meeting*, Volume 3, pp. 28182823.

Koesrindartoto, D. and L. Tesfatsion (2004). Testing the reliability of ferc's wholesale power market platform: an agentbased computational economics approach. In *Proceedings of the 24th USAEE/IAEE North American Conference, Washington, D.C., USA*.

Krause, A. (2000). *Market microstructure theory and strategic behavior of market makers*. Ph. D. thesis, University of Fribourg.

Krause, A. (2006). Fat tails and multi-scaling in a simple model of limit order markets. *Phisica A 368*, 183–190.

Krishna, V. (2002). *Auction Theory*. Academic Press.

Krishna, V. and V. Ramesh (1998). Intelligent agents for negotiations in market games, part 2, application. *IEEE Transactions on Power Systems 13*, 1103–1108.

Kumar, P. and D. Seppi (1994). Limit and market orders with optimising traders. Technical report, Working paper, Graduate School of Industrial Administration, Carnegie Mellon Univeristy.

Kunakote, T. and S. Bureerat (2008). Topology optimisation of a plate structure using multiobjective population-based incremental learning. In *Technology and Innovation for Sustainable Development Conference.*

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica 53*, 1315–1335.

Kyle, A. S. (1986a). Market structure, information, future markets and price formation. *International agricultural trade: advanced readings in price formation, market structure and price instability*, 45–64.

Kyle, A. S. (1986b). On the incentives to produce private information with continuous trade. *Working paper, University of California-Berkeley.*

Ladley, D. and K. R. Schenk-Hoppe (2007). Do stylised facts of order book markets need strategic behaviour? Swiss finance institute research paper series 07-20, Swiss Finance Institute.

Lamba, H. and T. Seaman (2006). Mean-square stability properties of an adaptive time-stepping sde solver. *Journal of Computational and Applied Mathematics 194*(2), 245–254.

Lane, D., A. Kroujiline, V. Petrov, and G. Sheble (2000). Electricity market power: marginal cost and relative capacity effects. In A. Alzala (Ed.), *Proceedings of the 2000 Congress on Evolutionary Computation (IEEE Society Press, Piscataway, N.J.)*, pp. 1048–054.

Larranaga, P., R. Etxeberria, J. A. Lozano, and J. M. Pena (1999). Optimization by learning and simulation of bayesian and gaussian networks. Technical report, Technical Report EHU-KZAA-IK-4/99, Department of Computer Science and Artificial Intelligence, University of Basque Country.

Larranaga, P. and J. Lozano (2001). *Estimation of Distribution Algorithms. A New Tool for Evolutionary Computation.* Boston, MA: Kluwer Academic.

Larranaga, P. and J. A. Lozano (2002). *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation.* Norwell, MA: Kluwer.

Lau, S. T. and T. H. McInish (1995). Reducing tick size on the stock exchange of singapore. *Pacific-Basin Finance Journal 3*, 485–496.

Laumanns, M., L. Thiele, K. Deb, and E. Zitzler (2002). Combining convergence and diversity in evolutionary multi-objective optimization. *Evolutionary Computation 10*, 263–282.

LeBaron, B. (2001). A builder's guide to agent based financial markets. *Quantitative Finance 1*, 254–261.

Lebaron, B. (2002, June). Building the santa fe artificial stock market. Working paper, Brandeis University.

LeBaron, B. (2005). *this handbook*, Chapter Agent-based computational finance.

LeBaron, B., W. B. Arthur, and R. Palmer (1999). Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control 23*, 1487–1516.

Lee, J.-Y., M.-S. Kim, and J.-J. Lee (2004). Multi-objective walking trajectories generation for a biped robot. In *Proceedings of 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems*, Volume 4, pp. 3853 – 3858.

Li, J. and E. P. K. Tsang (1999). Investment decision making using fgp: A case study. In *Proceedings of the Congress on Evolutionary Computation*, Volume 2, pp. 1253–1259.

Li, X. and A. Krause (2009a). A comparison of market structures with near-zero-intelligence traders. *Intelligent Data Engineering and Automated Learning - IDEAL 2009. Lecture Notes in Computer Science 5788/2009*, 703–710.

Li, X. and A. Krause (2009b). Determination of the optimal trading rules minimizing the bid-ask spread. University of Bath.

Li, X. and A. Krause (2010a). Determining the optimal market structure using near-zero intelligence traders. *Journal of Economic Interaction and Coordination 5*, 155–167. 10.1007/s11403-010-0069-3.

Li, X. and A. Krause (2010b). An evolutionary multi-objective optimization of market structures using pbil. *Intelligent Data Engineering and Automated Learning IDEAL 2010. Lecture Notes in Computer Science 6283/2010*, 78–85.

Littlechild, S. (1995). *Utility regulation: challenge and response*, Chapter Competition in electricity: restrospect and prospect, pp. 101–114. Institute of Economic Affairs.

Logue, D. E. (1975). Market-making and the assessment of market efficiency. *Journal of Finance 30*, 115–123.

Lovric, M. (2011). *Behavioral Finance and Agent-Based Artificial Markets*. Ph. D. thesis, Erasmus University Rotterdam.

Lucking-Reiley, D., D. Bryan, N. Prasad, and D. Reeves (1999). Pennies from ebay: the determinants of price in online auctions. Working paper, Vanderbilt University.

Lux, T. and M. Marchesi (2000). Volatility clustering in financial markets: a micro-simulation of interacting agents. *International Journal of Theoretical and Applied Finance 3*, 675–702.

Lyons, R. (1994). Optimal transparency in a dealership market with an application to foreign exchange. Working paper, University of California, Berkeley.

MacGill, I. and R. Kaye (1999). Decentralized coordination of power system operation using dual evolutionary programming. *IEEE Transactions on Power Systems 14*(1), 112–119.

Machado, P., J. Tavares, F. B. Pereira, and E. Costa (2002). Vehicle routing problem: Doing it the evolutionary way. In *In GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference.*

Madhavan, A. (1992). Trading mechanisms in securities markets. *The Journal of Finance 47*, 607–641.

Madhavan, A. (1996). Security prices and market transparency. *Journal of Financial Intermediation 5*, 255–283.

Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets 3*, 205–258.

Madhavan, A. and V. Panchapagesan (2000). Price discovery in auction markets: a look inside the black box. *The Review of Financial Studies 13*, 627–658.

Madhavan, A., D. Porter, and D. Weaver (2005). Should securities markets be transparent. *Journal of Financial Markets 8*, 265–287.

Madhavan, A. and S. Smidt (1993). An analysis of changes in specialist quotes and inventories. *Journal of Finance 48*, 1595–1628.

Madhavan, A. and G. Sofianos (1998). An empirical analysis of nyse specialist trading. *Journal of Finance 48*, 189210.

Maifeld, T. and G. Sheble (1996, Aug). Genetic-based unit commitment algorithm. *IEEE Transactions on Power Systems 11*(3), 1359–1370.

Major, D. (1969). Benefit-cost ratios for projects in multiple objective investment programs. *Water Resources Research 5*, 1174–1178.

Manaster, S. and S. C. Mann (1996). Life in the pits: Competitive market making and inventory control. *Review of Financial Studies 9*, 953–975.

Marglin, S. (1967). Public investment criteria. Technical report, MIT Press, Cambridge, MA.

Markowitz, H. (1952). The utility of wealth. *The Journal of Political Economy 60*(2), 151–158.

Marks, R. E. (2006). Market design using agent-based models. In L. Tesfatsion and K. L. Judd (Eds.), *Handbook of Computational Economics* (1 ed.), Volume 2, Chapter 27, pp. 1339–1380. Elsevier.

Marti, L., J. Garcia, A. Berlanga, and J. Molina (2009). An approach to stopping criteria for multi-objective optimization evolutionary algorithms: The mgbm criterion. In *IEEE Congress on Evolutionary Computation, 2009. CEC '09.*, pp. 1263 –1270.

Martinez-Jaramillo, S. and E. P. K. Tsang (2009). An heterogeneous, endogenous and co-evolution gp-based financial market. *IEEE Transaction on Evolutionary Computation 13*, 33–55.

Mas-Colell, A., M. D. Whinston, and J. R. Green (1995). *Microeconomic Theory*. Oxford University Press.

Maslov, S. (2000). Simple model of a limit order-driven market. *Physica A 278*, 571–578.

McAfee, R. P. and J. McMillan (1994). Selling spectrum rights. *Journal of Economic Perspectives 8*, 145–162.

McAfee, R. P. and J. McMillan (1996). Analyzing the airwaves auction. *Journal of Economic Perspectives 10*, 159–175.

Mendelson, M. (1972). From automated quotes to automated trading: Restructuring the stock market in the u.s. Technical report, Institute of Finance, New York University.

Menezes, F. M. and P. K. Monteiro (2008). *An Introduction to Auction Theory*. Oxford University Press.

Mhlenbein, H. and T. Mahnig (1999). The factorized distribution algorithm for additively decomposed functions. In *Proceedings of the 1999 Congress on Evolutionary Computation, IEEE press*.

Milgrom, P. R. (2004). *Putting auction theory to work*. Cambridge University Press.

Milgrom, P. R. and R. J. Weber (1982). A theory of auctions and competitive bidding. *Econometrica 50*(5), 1089–1123.

Mizuta, H. and K. Steiglitz (2000). Agent-based simulation of dynamic online auctions. In *In Proceedings of the 2000 Winter Simulation Conference*, pp. 1772–1777.

Mostaghim, S. and J. Teich (2004). Covering pareto-optimal fronts by subswarms in multi-objective particle swarm optimization. In *Congress on Evolutionary Computation, 2004. CEC2004.*, Volume 2, pp. 1404 – 1411.

Muhlenbein, H., J. Bendisch, and H.-M. Voigt (1996). From recombination of genes to the estimation of distributions. i. binary parameters. In *Proceedings of PPSN-I, First International Conference on Parallel Problem Solving from Nature*, pp. 178187. Berlin, Germany: Springer.

Myerson, R. B. (1981). Optimal auction design. *Mathematics of Operations Research 6*(1), 58–74.

Myerson, R. B. and M. A. Satterthwaite (1983). Efficient mechanisms for bilateral trading. *Journal of Economic Theory 29*(2), 265–281.

Naik, N., A. Neuberger, and S. Vishwanathan (1994). Disclosure regulation in competitive dealership markets: analysis of the london stock exchange. Working Paper, London Business School.

NASD, I. (1997). Market quality monitoring.

Nekovee, M. (2010). Mechanism design for cognitive radio networks. In *Proc. COMPENG '10. Complexity in Engineering*, pp. 12–17.

Nicolaisen, J., V. Petrov, and L. Tesfatsion (2001). Market power and efficiency in a computational electricitymarket with discriminatory double-auction pricing. In *IEEE Trans Evol Comput*, Volume 5, pp. 504–523.

Nicolaisen, J., M. Smith, V. Petrov, and L. Tesfatsion (2000). Concentration and capacity effects on electricity market power. In *in: A. Alzala, ed., Proceedings of the 2000 Congress on Evolutionary Computation*, pp. 1041–1047. IEEE Society Press, Piscataway, N.J.

Niu, J., K. Cai, S. Parsons, E. Gerding, P. McBurney, T. Moyaux, and S. P. D. (2008). Jcat: a platform for the tac market design competition. In *In: AAMAS '08: Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pp. 1649–1650.

Noe, T., M. J. Rebello, and J. Wang (2003). Corporate financing: An artificial agent-based analysis. *Journal of Finance 58*(3), 943–973.

Noe, T., M. J. Rebello, and J. Wang (2006). The evolution of security designs. *Journal of Finance 61*(5), 2103–2135.

Noia, C. D. (2001). Competition and integration among stock exchanges in europe: Network effects, implicit mergers and remote access. *European Financial Management 7*, 39–72.

Ockenfels, A. and A. Roth (2006). Late and multiple bidding in second price internet auctions: Theory and evidence concerning different rules for ending an auction. *Games and Economic Behavior 55*, 297–320.

O'Hara, M. (1995). *Market Microstructure Theory*. Blackwell, Oxford, Uk.

O'Hara, M. (2001). Designing markets for developing countries. *International Review of Finance 2:4*, 205–215.

O'Hara, M. (2007). Optimal microstructures. *European Financial Management 13*, 825–832.

O'Hara, M. and G. S. Oldfield (1986). The microeconomics of market making. *Journal of Financial and Quantitative Analysis 21*, 361–376.

O'Neill, M., A. Brabazon, C. Ryan, and J. J. Collins (2001). Evolving market index trading rules using grammatical evolution. In *In Proceedings of the EvoWorkshops on Applications of Evolutionary Computing*, pp. 343–352.

Ortega-Reichert, A. (1968). Models for competitive bidding under uncertainty. Technical report no. 8, Department of Operations Research, Stanford University.

Osyczka, A. (1985). Multicriteria optimization for engineering design. In I. J. G. editor (Ed.), *Design Optimization*, pp. 193C227. Academic Press.

Othman, A. (2008). Zero-intelligence agents in prediction markets. In *In: Proceedings of the Seventh International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2008)*, pp. 879–886.

Pagano, M., O. Randl, A. A. Roell, and J. Zechner (2001). What makes stock exchanges succeed? evidence from cross-listing decisions. *European Economic Review 45*, 770–782.

Pagano, M. and A. Roell (1996). Transparency and liquidity: A comparison of auction and dealer markets with informed trading. *The Journal of Finance 51*(2).

Pagano, M., A. Roell, and J. Zechner (2001). The geography of equity listing: Why do european companies list abroad? Technical Report 2681, C.E.P.R. Discussion Papers.

Pagano, M. and R. Schwartz (2003). A closing calls impact on market quality at euronext paris. *Journal of Financial Economics 68*, 439–484.

Panchapagesan, V. (1997). What if time priority is not enforced among traders? Working paper, University of Southern California.

Pareto, V. (1896). *Course of Political Economics*, Volume vol. 1 and 2. Lausanne:F. Rouge.

Parkes, D. C. (2008). Computational mechanism design. In *In Lecture notes of Tutorials at 10th Conf. on Theoretical Aspectsof Rationality and Knowledge (TARK-05)*. Institute of Mathematical Sciences, University of Singapore.

Parlour, C. A. (1998). Price dynamics in limit order markets. *The review of financial studies 11*, 789–816.

Pelikan, M., D. Goldberg, and E.Cantu-Paz (2000). Bayesian optimization algorithm, population sizing, and time to convergence. Technical report, University of Illinois at Urbana-Champaign, Illinois Genetic Algorithms Laboratory, Urbana, IL, IlliGAL Report No. 2000001.

Pelikan, M., D. Goldberg, and F. Lobo (1999). A survey of optimization by building and using probabilistic models. Technical report, Technical Report IlliGAL, 99018, University of Illinois at Urbana-Champaign, Urbana, IL.

Pelikan, M. and H. Muhlenbein (1999). The bivariate marginal distribution algorithm. *Advances in Soft Computing-Engineering Design and Manufacturing*, 521–535.

Peterson, R. L. (2007). *Inside the Investors Brain: The Power of Mind over Money.* John Wiley and Sons, Inc.

Petrov, V. and G. Sheble (2001). Building electric power auctions with improved roth-erev reinforced learning. In *Proceedings of the North American Power Symposium*, Texas, USA.

Phelps, S., P. McBurney, and S. Parsons (2010). Evolutionary mechanism design: a review. *Autonoumous Agents and Multi-Agent Systems 21*, 237–264.

Phelps, S., P. McBurney, S. Parsons, and E. Sklar (2002). Co-evolutionary auction mechanism design: A preliminary report. In *AAMAS '02: Revised Papers from the Workshop on Agent Mediated Electronic Commerce on Agent-Mediated Electronic Commerce IV, Designing Mechanisms and Systems*, London, UK, pp. 123–142. Springer-Verlag.

Phelps, S., S. Parsons, and P. McBurney (2005). An evolutionary game-theoretic comparison of two double-auction market designs. In P. Faratin and J. Rodriguez-Aguilar (Eds.), *Agent-Mediated Electronic Commerce VI*,

Volume 3435 of *Lecture Notes in Computer Science*, pp. 101–114. Springer Berlin / Heidelberg.

Phelps, S., S. Parsons, P. McBurney, and E. Sklar (2002). Co-evolution of auction mechanisms and trading strategies: Towards a novel approach to microeconomic design. *Proc. GECCO-02 Workshop Evolutionary Computation in Multi-Agent Systems*, 65–72.

Phelps, S., S. Parsons, E. Sklar, and P. McBurney (2003). Applying genetic programming to economic mechanism design: Ecolving pricing rule for continuous double auction. In J. S. Rosenschein, T. sandholm, M. Wooldridge, and M. Yokoo (Eds.), *Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-2003)*, pp. 1096–1097. Melbourbe, Australia.

Porter, D. and D. Weaver (1998a). Post-trade transparency on nasdaqs national market system. *Journal of Financial Economics 50*, 231–252.

Porter, D. and D. Weaver (1998b). Transparency and liquidity: Should u.s. markets be more transparent? Working paper, Marquette University.

Potvin, J.-Y., P. Soriano, and M. Vallee (2004). Generating trading rules on the stock markets with genetic programming. *Computers & Operations Research 31*(7), 1033–1047.

Preist, C. and M. van Tol (2003). Adaptive agents in a persistent shout double auction. Technical report, Technical Report HPL-2003-242, Hewlett Packard Labs.

Price, T. (1997). Using co-evolutionary programming to simulate strategic behavior in markets. *Journal of Evolutionary Economics 7*(3), 219–254.

Qin, Z. and T. Kovacs (2004). Evolution of realistic auctions. In *Proceedings of the 2004 UK workshop on Computational Intelligence*, Loughborough, UK, pp. 43–50.

Qin, Z. and T. Kovacs (2005). Evolutionary designs of mechanisms for markets with heterogeneous agents. In *Presented on 10th Annual Workshop on Economic Heterogeneous Interacting Agents (WEHIA 2005)*, University of Essex, UK.

Qin, Z. C. (2002). Evolving marketplace designs by artificial agents. Master's thesis, University of Bristol.

Qin, Z. C. (2006). Market mechanism designs with heterogeneous trading agents. In *Proceedings of ICMLA 2006: 5th International Conference on Machine Learning and Applications.*

Rabertoa, M., S. Cincotti, S. M. Focardi, and M. Marchesi (2001). Agent-based simulation of a financial market. *Physica A 299*, 319–327.

Ramos, S. B. (2003). Competition between stock exchanges: A survey. Technical report, FAME Research Paper No. 77.

Rasanen, M., R. Hamalainen, and J. Ruusunen (1994, Feb. 21?23). Visual interactive modeling in electricity load analysis. In M. Hamza (Ed.), *Modelling, Identification and Control, Proceedings the 13th International Association of Science and Technology for Development (IASTED) International Conference*, Grindelwald, Switzerland, pp. 339–342.

Rastegar, R. and M. R. Meybodi (2005). A new estimation of distribution algorithm based on learning automata. *Proceedings of the IEEE 2005 Congress on Evolutionary Computation*, 1982–1987.

Rhee, S., G. Stone, and N. Tang (2004). The impact of trading method on volatility and market manipulation: New evidence from the singapore exchange. University of Hawaii Working Paper.

Richardson, J. T., M. R. Palmer, G. E. Liepins, and M. R. Hilliard (1989). Some guidelines for genetic algorithms with penalty functions. In *Proceedings of the third international conference on Genetic algorithms*, pp. 191–197.

Richter, C., G. S. Jr., and D. Ashlock (1999, November). Comprehensive bidding strategies with genetic programming/finite state automata. *IEEE Transactions on Power Systems 14*(4), 1207 –1212.

Ricker, J. P. (1999). Breaking the eighth: sixteenths on the new york stock exchange.

Riley, J. G. and W. F. Samuelson (1981). Optimal auctions. *The American Economic Review 71*(3), 381–393.

Ritter, J. R. (2003). Behavioral finance. *Pacific-Basin Finance Journal 11*, 429–437.

Robinson, N. (2002). Evolutionary optimization of market-based control systems for resource allocation in compute farms. Master's thesis, School of Cognitive and Computing Sciences, University of Sussex, U.K.

Rock, K. (1988). The specialist's order book: A possible explanation for the year-end anomaly. Working paper, Harvard Business School.

Roell, A. (1990). Dual capacity trading and the quality of the market. *Journal of Financial Intermediation 1*, 105–124.

Roger, P. and L. Eeckhoudt (1999). Risk aversion and the bid-ask spread. *European Financial Management 5*, 323–340.

Ronen, T. and D. G. Weaver (1998). The effect of tick size on volatility, trader behavior, and market quality. Working paper, Baruch College, New York.

Rosenberg, R. (1967). *Simulation of genetic population with biochemical properties.* Phd thesis, University of Michigan.

Roth, A. (2002). The economist as engineer: Game theory, experimental economics and computation as tools of design economics. *Econometrica 70*, 1341–1378.

Roth, A. and A. Ockenfels (2002). Last minute bidding and the rules for ending second-price auctions: evidence from ebay and amazon auctions on the internet. *American Economic Review 92*, 1093–1103.

Rudlof, S. and M. Koppen (1996). Stochastic hill climbing with learning by vectors of normal distributions. In *in First On-line Workshop on Soft Computing, Nagoya, Japan*.

Rudolph, G. (1998). Evolutionary search for minimal elements in partially ordered finite sets. In *in Proceedings of the 7th Annual Conference on Evolutionary Programming*, pp. 345–353. Springer, Berlin.

Rudolph, G. (1999). Evolutionary search under partially ordered sets. In *in Proceedings of the International NAISO Congress on Information Science Innovations*.

Rudolph, G. and A. Agapie (2000). Convergence properties of some multi-objective evolutionary algorithms. In *in Proc. of IEEE Congress on Evolutionary Computation*.

Rust, J. and G. Hall (2003). Middlemen versus market makers: a theory of competitive exchange. *Journal of Political Economy 111*, 353–403.

Saporta, V., G. Trebeschi, and A. Vila (1999). Price formation and transparency on the london stock exchange. Working Paper 95, Bank of England.

Sarker, R. A. and T. Ray (2010). *Agent-based Evolutionary Search*. Springer-Verlag Berlin Heidelberg.

Sawaragi, Y., H. Nakayama, and T. Tanino (1985). *Theory of multiobjective optimization*. Orlando, Florida: Academic Press.

Schaffer, D. J. (1985). Multiple objective optimization with vector evaluated genetic algorithms. In *Proceedings of the 1st International Conference on Genetic Algorithms*, pp. 93–100. L. Erlbaum Associates Inc.

Schaffer, J. D. and A. Morishima (1987). An adaptive crossover distribution mechanism for genetic algorithms. In *Proceedings of the Second International Conference on Genetic Algorithms*, pp. 36–40.

Schindler, M. (2007). *Rumors in financial markets: insights into behavioral finance.* John Wiley & Sons, Ltd.

Schroder, P., A. Chipperfield, P. J. Fleming, and N. Grurn (1997). Multiobjective optimizationof distributed active magnetic bearing controllers. In *In: Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications*, pp. 13–18.

Schulenburg, S. and P. Ross (2000). An adaptive based economic model. *Lecture Notes in Artificial Intelligence*, 263–282.

Schwartz, R. A. (1988). *Equity markets: structure, trading, and performance.* New York: Harper & Row.

Sebag, M. and A. Ducoulombier (1998). Extending population-based incremental learning to continuous search spaces. In *in Parallel Problem Solving from NaturePPSN V*, pp. 418–427. Berlin Heidelberg.

Seifert, S. and K.-M. Ehrhart (2005). Design of the 3g spectrum auctions in the uk and germany: An experimental investigation. *German Economic Review 6*, 229–248.

Seppi, D. J. (1997). Liquidity provision with limit orders and a strategic specialist. *Review of Financial Studies 10*, 103–150.

Servet, I., L. Trave-Massuyes, and D. Stern (1997). Telephone network traffic overloading diagnosis and evolutionary computation techniques. In *In Third European Conference on Artificial Evolution (AE'97)*, pp. 137–144. Berlin. Springer Verlag.

Sewell, M. (2011). Characterization of financial time series. Research Note RN/11/01, University College London, London.

Shefrin, H. (2000). *Beyond Greed and Fear - Understanding Behavioral Finance and the Psychology of Investing.* Harvard Business School Press.

Shimokawa, T., K. Suzuki, and T. Misawa (2007). An agent-based approach to financial stylized facts. *Physica A 379*, 207–225.

Shleifer, A. and L. H. Summers (1990). The noise trader approach to finance. *Journal of Economic Perspectives 4*, 19–33.

Slovic, P. (1969). Analyzing the expert judge: A study of a stockbroker's decision process. *Journal of Applied Psychology 53:1*, 255–263.

Smith, V. (1962). An experimental study of competitive market behavior. *Journal of Political Economy 70*, 111–137.

Southey, F. and F. Karray (1999). Approaching evolutionary robotics through population-based incremental learning. In *In Proceedings of the 1999 International Conference on Systems, Man, and Cybernetics (SMC'99).*

Srinivas, N. and K. Deb (1994). Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation 2*, 221–248.

Steuer, R. E. (1986). *Multiple criteria optimization: Theory, computations, and application.* New York: John Wiley and Sons, Inc.

Stoll, H. R. (1978). The supply of dealer services in securities markets. *Journal of Finance 33*, 1133–1151.

Sun, J. and L. Tesfatsion (2007). Dynamic testing of wholesale power market designs: an open-source agent- based framework. *Computational Economics 30*(3), 291–327.

Talukdar, S. (2002, December 10). Agent-based market testing. In *DOE Transmission Reliability Research Review.* Washington DC.

Talukdara, S. and V. Ramesh (1992). A-teams for real-time operations. *International Journal of Electrical Power and Energy Systems 14*(2-3), 138–143.

Tan, K. C., E. F. Khor, and T. H. Lee (2005). *Multiobjective Evolutionary Algorithms and Applications.* Springer.

Tan, K. C. and Y. Li (1997). Multi-objective genetic algorithm based time and frequency domain design unification of linear control systems. Technical report, Department of Electronics and Electrical Engineering, University of Glasgow.

Tesfatsion, L. (2001). Introduction to the special issue on agent-based computational economics. *Journal of Economic Dynamics and Control 25*(3-4), 281–293.

Tesfatsion, L. (2002). Agent-based computational economics: Growing economies from the bottom up. *Artificial Life 8*, 55–82.

Tesfatsion, L. (2003). Agent-based computational economics: modeling economies as complex adaptive systems. *Information Sciences 149*(4), 262–268.

Tesfatsion, L. (2006). Agent-based computational economics: A constructive approach to economic theory. In L. Tesfatsion and K. Judd (Eds.), *Handbook of Computational Economics*, Volume 2 of *Handbook of Computational Economics*, pp. 831 – 880. Elsevier.

Thaler, R. H. (Ed.) (2005). *Advances in Behavioral Finance.* Princeton University Press.

Theissen, E. (2003). Trader anonymity, price formation and liquidity. *European Finance Review 7*, 1–26.

Thompson, H. and P. Fleming (1999). An integrated multidisciplinary optimisation environment for distributed aero-engine control system architectures. In *In: Proceedings 14th World Congress of International Federation of Automatic Control*, pp. 407–412.

van Damme, E. (2002). The european umts-auctions. *European Economic Review 46*, 846–858.

Van Ness, B. F., R. A. V. Ness, and S. W. Pruitt (2000, September). The impact of the reduction in tick increments in major u.s. markets on spreads, depth, and volatility. *Review of Quantitative Finance and Accounting 15*(2), 153–67.

Van Veldhuizen, D. A. (1999). *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations.* Ph. D. thesis, Graduate School of Engineering of the Air Force Institute of Technology, Air University.

Varian, H. R. (1995). Economic mechanism design for computerized agents. In *In USENIX workshop on electronic commerce, New York, NY.*

Veit, D., A. Weidlich, J. Yao, and S. Oren (2006). Simulating the dynamics in two-settlement electricity markets via an agent-based approach. *International Journal of Management Science and Engineering Management 1*, 83–97.

von der Fehr, N.-H. M. and D. Harbord (1993). Spot market competition in the uk electricity industry. *The Economic Journal 103*(418), 531–546.

von der Fehr, N.-H. M. and D. Harbord (1998). Competition in electricity spot markets. economic theory and international experience. Memorandum 05/1998, Oslo University, Department of Economics.

Waisburd, A. (2003). Anonymity and liquidity: evidence from the paris bourse. mimeo,Texas Christian University.

Walia, V. (2002). Evolving market design. Technical Report HPL-2002-289, HP Laboratories Bristol.

Wang, Z., T. Chen, K. Tang, and X. Yao (2009). A multi-objective approach to redundancy allocation problem in parallel-series systems. In *IEEE Congress on Evolutionary Computation, 2009. CEC '09.*, pp. 582 –589.

Weaver, D. G. (2006). Networks, nodes, and priority rules. In C.-F. Lee and A. C. Lee (Eds.), *Encyclopedia of Finance*, pp. 689–699. Springer US.

Weber, B. W. (2006). Adoption of electronic trading at the international securities exchange. *Decision Support Systems 41*, 728–746.

Weidlich, A. and D. Veit (2008, July). A critical survey of agent-based wholesale electricity market models. *Energy Economics 30*(4), 1728–1759.

Westerhoff, F. (2003). Market-maker, inventory control and foreign exchange dynamics. *Quantitative Finance 3*, 363–369.

Wilcox, R. T. (2000). Experts and amateurs: The role of experience in internet auctions. *Markeing Letters 11*(4), 363–374.

Wilson, R. (2002). Architecture of power markets. *Econometrica 70*(4), 1299–1340.

Wilson, R. B. (1967). Competitive bidding with asymmetric information. *Management Science 13*(11), 816–821.

Wilson, R. B. (1969). Competitive bidding with disparate information. *Management Science 15*(7), 446–448.

Wintr, L. (2004). Some evidence on late bidding in ebay auctions. Working paper, Clark University, Worcester MA.

Yang, S. and X. Yao (2003). Dual population-based incremental learning for problem optimization in dynamic environments. In *Proceedings of the 7th Asia Pacific Symposium on Intelligent and Evolutionary Systems,*, pp. 49–56.

Yang, S. and X. Yao (2005). Experimental study on population-based incremental learning algorithms for dynamic optimization problems. *Soft Computing 9*, 815–834.

Yeh, C. H. (2003). Tick size and market performance. Technical report, Technical report, Yuan Ze University, Chungli, Taiwan.

Yuret, D. and M. Maza (1994). A genetic algorithm system for predicting the oex. *Technical Analysis of Stocks & Commodities*, 58–64.

Zhang, Q. (2003). A rigorous investigation into estimation of distribution algorithms. Technical report, IGR Report for EPSRC Grant GR/R64742/01.

Zhang, Q. and H. Li (2007). Moea/d: A multi-objective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation 11*, 712–731.

Zitzler, E. (1999). *Evolutionary algorithms for multiobjective optimization: methods and applications.* Ph. D. thesis, Swiss Federal Institute of Technology Zurich.

Zitzler, E., K. Deb, and L. Thiele (2000). Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation 8*, 173–195.

Zitzler, E., M. Laumanns, and L. Thiele (2001). Spea2: Improving the strength pareto evolutionary algorithm. In K. C. Giannakoglou, D. T. Tsahalis, J. Périaux, K. D. Papailiou, and T. Fogarty (Eds.), *Evolutionary Methods for Design Optimization and Control with Applications to Industrial Problems*, pp. 95–100. International Center for Numerical Methods in Engineering.

Zitzler, E. and L. Thiele (1999). Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation 3(4)*, 257–271.