QUANTITATIVE FINANCE
RESEARCH CENTRE

UNIVERSITY OF
TECHNOLOGY SYDNEY

THINK.CHANGE.DO

QUANTITATIVE FINANCE RESEARCH CENTRE

Research Paper 335 August 2013

Learning and Evolution of Trading Strategies in Limit Order Markets

Carl Chiarella, Xue-Zhong He and Lijian Wei

ISSN 1441-8010 www.qfrc.uts.edu.au

LEARNING AND EVOLUTION OF TRADING STRATEGIES IN LIMIT ORDER MARKETS

CARL CHIARELLA, XUE-ZHONG HE AND LIJIAN WEI*

UTS Business School, University of Technology, Sydney
PO Box 123, Broadway, NSW 2007, Australia

Abstract. How do traders process and learn from market information, what trading strategies should they use, and how does learning affect the market? This paper proposes a learning model of an artificial limit order market with asymmetric information to address these issues. Using a genetic algorithm as a learning mechanism, we show that learning, in particular the learning from uninformed traders, improves market informational efficiency and has a significant impact on the stylized facts of limit order markets, order submission, liquidity supply and consumption, the hump shaped order book near the quote, and the bid-ask spread. Moreover, the learning affects the evolution process of the trading strategies for all traders. The model provides some insights into market efficiency, the interaction of traders, the dynamics of limit order books, and the evolution of trading strategies.

Key words: Limit order book, evolution, genetic algorithm learning, asymmetric information, trading strategy.

JEL Classification: G14, C63, D82

Date: July 15, 2013.

Acknowledgement: We thank the participants at the Workshop on "Market Microstructure and Nonlinear Dynamics" (MMND), the University of Evry (France) on June 13-14, 2013, in particular, Bruce Lehmann for helpful comments. We thank our colleagues Anthony Hall, Talis Putnins, Lei Shi and Jianxin Wang for valuable suggestions. Lijian Wei thanks Wei Zhang and Yongjie Zhang for the initial discussion of the two-sided learning mechanism with the genetic algorithm during his PhD studies at Tianjin University. Financial support from the Australian Research Council (ARC) under a Discovery Grant (DP110104487) is gratefully acknowledged.

*Corresponding author.

Emails: `Carl.Chiarella@uts.edu.au (Chiarella), Tony.He-1@uts.edu.au(He), weilijian@gmail.com (Wei)`.

1. Introduction

Limit order markets are the most popular and efficient financial markets, in which continuous double auction is one of the most important mechanisms. It provides trading opportunities for different types of traders with various trading strategies continuously and helps them to learn and to choose the best strategies when market conditions are changed. A trading strategy in limit order markets may include buy or sell, order type (limit or market orders), order size, time to trade, time to re-enter, order cancelation, etc. For example, traders may use market orders if they believe the current quote is far away from the fundamental value, and use limit orders away from the quote when the quote is close to the fundamental value; when the bid-ask spread is wide, they may use aggressive limit orders inside the bid-ask spread to replace market orders to improve order profit with a high trading priority. Therefore trading strategies affect market information efficiency and order book dynamics, including order submission, order book shape, and bid-ask spread, which in turn affect the evolution of trading strategies of traders through learning. Hence, the modeling of trading strategies and learning to capture the limit order market dynamics is very important.

Due to the complexity of limit order markets, it is very challenging to combine all the decisions into trading strategies and to model how trading strategies are formed and evolve dynamically. Most of the models in the literature focuses on one or two decisions, in particular, the buy or sell decision and the choice of order type. Some early static models of limit order markets assume that informed traders only use market orders and uninformed traders or liquidity traders only use limit orders (see, for example, Glosten (1994) and Seppi (1994)). However empirical and theoretical studies find that both informed and uninformed traders use mixed market orders and limit orders.[1] For more discussion about the choice of order type, we refer the reader to Parlour and Seppi (2008) and Rosu (2012*b*). Goettler, Parlour and Rajan (2009) and Rosu (2012*a*) are the two typically dynamic models that allow informed and uninformed traders to determine their choice of order type in limit order markets.

Goettler et al. (2009) develop the first dynamic model with asymmetric information in limit order markets. It assumes that the information is short-lived, meaning that informed traders know the current fundamental value while uninformed traders know the fundamental value with a time lag. With exogenously given private values, which are measured by the deviations from the fundamental value, traders with high private value are uninformed traders and traders with zero private value are

[1]See, for example, Keim and Madhavan (1995), Chakravarty and Holden (1995), Harris (1998), Bloomfield, O'Hara and Saar (2005), Kaniel and Liu (2006), Menkhoff, Osler and Schmeling (2010) and Collin-Dufresne and Fos (2013).

informed traders. Due to an information acquisition cost, traders choose buy or sell and order type by a trade-off among the information cost, private value, the expected order profit (the difference between expected transaction price and fundamental value), and pick-off risk. Hence trading strategies of the uninformed traders are mainly determined exogenously by private value; those with positive (negative) private values prefer to buy (sell) and submit market or limit orders, depending on how high (low) the private values are. Informed traders with zero private value prefer to buy or sell with either market or limit orders, depending on their information advantage and the pick-off risk. Therefore, using limit orders to provide liquidity is profitable for the informed traders since the uninformed traders consume liquidity. It is found that in general, the informed traders are speculators who prefer to submit limit orders and supply liquidity. However, when the volatility of the fundamental value is high, the informed traders prefer to submit more market orders to reduce pick-off risk and profit from mispriced orders in the limit order book.

Similar to Goettler et al. (2009), Rosu (2012*a*) builds a dynamical model in which the buy/sell decisions of uninformed traders are also exogenously determined. But the order-choice of the uninformed traders depends on their time preference, meaning that traders are either patient or impatient. The patient traders prefer to use limit orders and the impatient traders prefer to use market orders. The model has two main differences from Goettler et al. (2009). The first one is that the information is long-lived, meaning that the informed traders with zero private value pay a fixed cost and observe the fundamental value when they enter the market, but the precision of their private information decays over time. The second difference is that traders can continuously monitor the market and cancel the limit order at any time so that the pick-off risk for the informed traders is eliminated. Because of the time preference, all traders on the same side of the book have the same expected utility. Rosu finds that in equilibrium, the patient informed traders submit both limit and market orders, depending on their information advantage, while the patient uninformed traders submit limit orders. However the impatient traders always submit market orders.

The two dynamic models contribute significantly to our understanding of limit order markets. However, they mainly focus on the trading strategies of informed traders, while the behavior of uninformed traders is simplified. In particular, the results highly depend on exogenous parameters, such as private value and time preference. The models do not consider how traders learn from market information and do not allow traders to choose trading strategies endogenously based on market conditions. As pointed out by O'Hara (2001),"*It is the uninformed traders who provide the liquidity to the informed, and so understanding their behaviors can* *provide substantial insight and intuition into the trading process. Information-based microstructure models typically assume that uninformed traders do not act strategically. Yet, if it is profitable for informed traders to time their trades, then it must be profitable for uninformed traders to do so as well.*" She further highlights the importance of learning, "*Another open question is what traders can learn from other pieces of market data, such as prices...Technical analysis of market data is widespread in markets, with elaborate trading strategies devised to respond to the pattern of prices.*" This paper aims to address these challenges, in particular on how traders learn and choose their trading strategies based on market conditions and how they interact with each other in a limit order market.

In this paper, we use a genetic algorithm (GA) mechanism to model traders' learning. The GA is a search heuristic that mimics the process of natural evolution. It generates solutions to optimization problems using techniques inspired by natural evolution, such as mutation, selection, and crossover. Apart from allowing agents to learn from experience, the GA is spontaneous and creative (Chen, Chang and Du (2012)). In the initial stage, the GA randomly generates some solutions, and evaluates them by their performance. Then, the GA uses three processes including selection, crossover and mutation to evolve the solutions based on the survival of the fittest, and creates new solutions from solutions with good historical performance and uses them to replace solutions with bad historical performance. In this way, the GA generates optimal sets of solutions to fit changes in the environment.

Since introduced first by Holland (1975), the GA has been widely used in economics and finance as an adaptive way to model investor's learning behavior.[2] A typical application of the GA to financial markets is the Santa Fe Institute Artificial Stock Market (SFI-ASM, Arthur, Holland, LeBaron, Palmer and Tayler (1997)). The SFI-ASM allows traders to forecast price by developing a classifier system based on some classify rules on market conditions. Classify rules, such as *"the current ask is higher than the fundamental value"*, describe the market conditions. A trading strategy contains two parts: market conditions and actions. The market conditions may include market information of bid, ask, bid-ask midpoint, market price, historical prices, and order book depthetc. The actions may include buy/sell, market order, aggressive limit order, limit order at bid/ask, and unaggressive limit order etc. For example, one trading strategy can be: *"when the ask is higher than the expected fundamental value and the current bid-ask spread is lower than before, then traders choose market buy"*. Trading strategies are then evolved according to their historical performance. When a trader enters the market, he/she chooses the best strategies from selected candidates that match the current market conditions.

[2]See, for example, We refer readers to Holland and Miller (1991), Arifovic (1994, 1996), Routledge (1999, 2001), and Chen (2002).

Arthur et al. (1997) show that the adaptive learning makes market price converge to the rationally expected price. Neo et al. (2003, 2006) apply the GA learning to model corporate financing and reveal the evolution process of security design. To generate some intraday trading patterns in limit order markets, Kluger and McBride (2011) allow both informed and uninformed traders using GA to decide when to enter the market during a trading day. More recently, Wei, Zhang, He and Zhang (2013*a*, 2013*b*) build a learning model of the limit order book under a similar information structure to Goettler et al. (2009). It allows part of the uninformed traders to learn to use the GA. The model is able to generate some stylized facts, such as fat tails, volatility clustering and long memory, and limit order phenomena. They show that learning of the uninformed traders improves information dissemination efficiency. More important, it provides a method on how to use a classifier system to describe the market conditions, so that traders can optimize their forecasting rules based on market information.

Motivated by Wei et al. (2013*a*, 2013*b*), this paper proposes a learning and evolution model of a limit order market that allows both informed and uninformed traders to learn so that their trading strategies evolve endogenously according to market conditions. We extend the classifier system used in Wei et al. (2013*a*, 2013*b*) and introduce more detailed classify rules to deal with more complicated market conditions. To compare with the two dynamic models, in particular Goettler et al. (2009), we assume a similar information structure and fundamental value process. However, unlike Goettler et al. (2009), the trading strategies do not depend on the exogenously given private value and time preference. Instead, they are determined by the GA based on the private information and market conditions. Also, different from Rosu (2012*a*), to facilitate the learning, traders are allowed to re-enter the market following a Poisson process and cancel any pending limit orders, but they do not always monitor the market and can not freely cancel an order at any time[3]. In particular, the learning model introduced in this paper aims to answer three important questions: (i) can learning and evolution of trading strategies improve market informational efficiency? (ii) how does the evolution of trading strategies affect the dynamics of the order book and traders' order submission behavior? (iii) how do the informed and uninformed traders process the private and market information and interact with each others? The answers to these questions can help us to understand how traders make their decisions on order submission, reveal the formation mechanism of limit order book phenomena, and provide valuable implications for algorithmic trading that has been used widely in limit order markets. As far as we

[3]This is similar to Dugast (2012) who proposes a model based on Rosu (2012*a*) where agents have limited attention and thus do not monitor the market continuously.

are aware, this is the first paper that uses GA learning to model fully endogenized trading strategies in limit order markets.

We now summarize the main findings and discuss the main contributions of the paper. To help the discussion, we refer *one-sided* learning to the cases where either the informed or the uninformed traders learn and *two-sided* learning to the case where both the informed and uninformed traders learn.

(i) Learning, in particular the learning from the uninformed traders, becomes very effective. It helps the traders to develop and update their trading strategies and makes the market price converge to the fundamental value. Most importantly, learning from the uninformed traders improves market informational efficiency and contributes to the stylized facts in limit order markets. The effect of learning is consistent with Goettler et al. (2009) and Rosu (2012*a*) for informed traders and Wei et al. (2013*a*,2013*b*) for uninformed traders.

(ii) Learning effects order submission significantly. When both the informed and uninformed traders learn, the informed traders increase market orders and reduce aggressive limit orders, comparing to the uninformed traders. Therefore the informed traders are more willing to consume liquidity than the uninformed traders. This result is different from Goettler et al. (2009) and Rosu (2012*a*). Comparing to Goettler et al. (2009), the uninformed traders do not always prefer to use market orders and the limit orders from the informed traders become less profitable. Instead of reducing the pick-off risk and obtaining order profit from the short-lived information, the informed traders prefer to use more market orders to consume liquidity, while the uninformed traders prefer to use more limit orders to supply liquidity. This contribution is new to the literature. Consistent with Goettler et al. (2009), when the fundamental volatility becomes higher, the informed traders submit more market orders.

(iii) Compared to one-sided learning, under two-sided learning, the uninformed traders use market orders or more limit orders at the quote to replace some aggressive limit orders, thus consume more liquidity. Menkhoff et al. (2010) find that the uninformed traders generally treat aggressive limit orders as an alternative to patient limit orders. Furthermore, they find that the placement of aggressive limit orders by the informed is fairly sensitive to changing market conditions, but not so for the uninformed traders. Our results show that when the uninformed traders learn, they are also sensitive to the changing market conditions.

(iv) The model is able to generate a "hump" shaped order book near the quote and the learning reduces order imbalance for quotes beyond the bid and ask. We show that, without learning but with rich order types, the continuous double auction generates a "hump" shape near the quote in the order book. Rosu (2009) conjectures that patient traders optimistically consider limit orders that are far away from the bid and ask which might be executed in the future, while impatient traders pessimistically believe that these order will never be matched and executed. It thus generates the "hump" near the quote and beyond the quote. Our result provides an explanation from the perspective of market mechanism rather than focusing on traders' behaviors. As a supplement to Rosu (2009), we show that limit order book depth, in particular the imbalance of depth beyond the bid and the ask contains valuable information, so that when traders learn, they can reduce the imbalance.

(v) Learning has a different impact for the informed and uninformed traders on the evolution of their trading strategies. For the informed traders, they pay more attention to market conditions than the uninformed traders. In particular, they care more about the change in market price and the depth at the quote. This is consistent with the finding in Menkhoff et al. (2010) that informed traders are more sensitive to the changing market conditions than uninformed traders. The asymmetric sensitivity is due to the asymmetric information. The learning from the uninformed traders increases the using frequency of all the classified rules (CRs) for the informed traders, in particular the CRs related to the price trend and the sign of the last transaction. When the market becomes less informative, the informed traders care more about the imbalance between the current depth at the ask and the current depth at the bid but less about other depth changes.

(vi) For the uninformed traders, they care mostly about the market conditions. This is different from Menkhoff et al. (2010) who find that uninformed traders are relatively muted in response to many changing market conditions, which are less intuitive. Furthermore, learning from the informed traders does not affect significantly the use of the CRs for the uninformed traders, although they care more about the relation between the bid-ask midpoint and the fundamental value, and the imbalance of the depth beyond the quote and the imbalance of the depth of the buy side and the sell side, but less about the market price and the historical price. When the market becomes more informative, the uninformed traders care more about the relation between the market price and the historical price and all the information related to the limit order book. Furthermore, when less weight is associated with the latest performance, the uninformed traders care less about the market conditions. When information-lived time is shorter, the uninformed traders care more about the market conditions.

The last two findings help us to understand how the traders process information and respond to the change in market conditions, and how they interact with each other. These results are new to the literature and provide valuable implications for the algorithmic trading in financial markets.

The rest of paper is organized as follows. The model is outlined in Section 2. Section 3 examines the impact of learning on price convergence, informational efficiency and stylized facts. Section 4 focuses on order submission behavior and the impact on order book shape, bid-ask spread and order profit. Section 5 examines the evolution process of trading strategies and the interaction of informed and uninformed traders. Section 6 concludes.

2. The Model

We consider an artificial limit order market which employs a continuous double auction trading mechanism. Traders are either informed or uninformed and their trading strategies are generated and updated endogenously through genetic algorithm learning based on private and public information.

2.1. **The limit order market.** There are N risk neutral traders and each trader arrives at the market according to a Poisson process with parameter λ. We assume that there are N_I informed and N_U uninformed traders so that $N_I + N_U = N$. The informed traders know the fundamental values when they arrive at the market, but not for the uninformed traders. The information is short-lived, meaning that the uninformed traders know the fundamental values with a time lag $\tau > 0$, we also call τ the information-lived time. The information structure is the same as in Goettler et al. (2009). A trading time period t, defined by $(t-1, t]$, corresponds to a short time interval in the real market, for instance one minute would be a typical value. The fundamental value v_t of the risky asset at time period t follows a random walk process. Innovations in the fundamental value v_t occur according to a Poisson process with parameter ϕ and initial fundamental value v_o. If an innovation occurs, the fundamental value either increases or decreases with equal probability by κ tick sizes. Depending on the value of the parameter ϕ, an innovation may occur in more than one time period when $\phi < 1$. All the informed traders who enter the market in time period t know the (same) fundamental value v_t; however the uninformed trader knows the fundamental value $v_{t-\tau}$ and $\tau > 0$ is measured in units of a time period. In general, the time lag τ can be different for different uninformed traders. In this paper, for simplicity we keep the time lag the same for all the uninformed traders; however we vary τ from 30 up to 120 time periods (namely half an hour up to 2 hours if one minute is taken as the time unit) in different scenarios to examine the effect of the information lag τ. Transactions take place based on the standard price and time priorities in limit order markets. When trader j enters the market at time $t' \in (t-1, t]$ in time period t, he observes a number of pieces of common information from the market price and the limit order book, including the current transaction price $p_{t'}$, the buy or sell initiated transaction sign $p^{\pm}_{t'}$ (+ for a buy and - for a sell), the most recent historical price $p_{t'-1}$, the average market price $\bar{p}_{t,\tau} = [p_{t-1} + p_{t-2} + \cdots + p_{t-\tau}]/\tau$ over the last τ periods, the mid-price

$p^m_{t'} = (a_{t'} + b_{t'})/2$ of the current bid $(b_{t'})$ and ask $(a_{t'})$ prices, the current bid-ask spread $s_{t'} = a_{t'} - b_{t'}$, the depth of the limit order book, the depth at the bid $d^b_{t'}$ and the ask $d^a_{t'}$, the depth below the bid $d^{bb}_{t'}$ and the ask $d^{aa}_{t'}$. We let $p_t = p_{t-1}$ if there is no transaction between time $t-1$ and t. To simplify the limit order book in extreme market conditions, we let $a_{t'} = 1.01 p_{t'}$ when the sell limit order book is empty and $b_{t'} = 0.99 p_{t'}$ when the buy limit order book is empty.[4]. The limit order expires in time T, which is set to be one day. We allow traders to reenter the market and cancel the previous order and submit a new order upon reentry.

2.2. **Trading strategies with a genetic algorithm.** When a trader enters the market, he uses a genetic algorithm (GA) to learn from market information on the market price and the limit order book and chooses the best trading strategy to buy or sell one share[5] with either a market order, or an aggressive limit order, or a limit order at quote, or an unaggressive limit order. The only difference between the informed and uninformed traders is that the trading decision to buy or sell is determined by the private information of the fundamental value for the informed traders but it is part of learning for the uninformed traders.

GA learning is based on the principles of natural selection. The outcome or solution of GA learning is called a chromosome, which is evaluated based on the its historical performance and selectively evolved through processes of selection, crossover and mutation (to be defined later). In the framework of the Santa Fe Institute Artificial Stock Market (SFI-ASM), a classifier system is introduced so that an agent can recognize market conditions and choose the chromosome accordingly. For our

[4]Actually, there is no bid when the buy limit order book is empty or no ask when the sell limit order book is empty. However, our model needs to provide the bid and the ask to the traders so that they can generate their trading strategies. This is a simple but reasonable method to solve the extreme limit order book condition. For example, if the transaction price $p_{t'} = 20$, and only when the sell limit order book is empty, and the bid $b_{t'} \leq 20$, we set the current ask $a_{t'}$ to 20.20, since the tick size is set to 0.02, the bid-ask spread is 10 tick sizes. If the sell limit order book keeps empty and only aggressive limit buy orders arrive at the market, $p_{t'}$ and the ask keep the same, but the bid goes up so that it may cross the ask, if this case happens, once the sell limit order with a price equal or lower than the bid arrives at the market, we let the cross order be executed immediately. In real limit order markets, traders often use this type limit orders with cross prices, which are called marketable limit orders.

[5] As pointed out by Rosu (2012*b*), most of limit order market models assumes risk-neutral traders with order size of one. Considering the order size decisions is important, which has be partially examined by an agent-based model of Chiarella, Iori and Perellò (2009) with exogenous order submission rule under symmetric information. To endogenize order size decision and order-choice with asymmetric information is an important but complicated issue. We also conduct a test with random order size between 1 to 10 and find that the results are the same. We conjecture that, if the traders are allowed to endogenously determine order size, the results would be different, in particular in the information efficiency, order book shape and the information processing of traders.

model, a chromosome corresponds to the trading strategy of a trader. Following the SFI-ASM, we use the classifier system to describe market conditions of the limit order books. A trading strategy i contains two components. The first component is called the use condition x^i (e.g. the current trading price $p_{t'}$ is larger than the average market price $\bar{p}_{t,\tau}$), which corresponds to market conditions. When x^i satisfies the current market condition at t', the trading strategy i is selected and added to an active candidate list. The trader then chooses the best trading strategy i^* from the list according to its strength, which is mainly determined by its historical performance (to be specified later). The second component is an action (order type) y^i of buying or selling and order aggressiveness. Once the best trading strategy i^* has been chosen, the trader uses the action y^{i^*} to trade. We now provide some details on the two components of a trading strategy based on GA learning.

The use condition x^i is based on the classification of market conditions. The classified rules (CRs) of the classifier system in the GA for our model are motivated by Goettler et al. (2009), Menkhoff et al. (2010) and Wei et al.(2013*a*, 2013*b*). Goettler et al. (2009) find that the change in ask/bid, the last transaction price, the last transaction sign (buy or sell), depth at ask/bid, and depth above the ask and below bid significantly affect the expected fundamental value for the uninformed traders. Menkhoff et al. (2010) find that the order submission of both informed and uninformed traders are sensitive to the spreads, volatility, momentum and depth. When uninformed traders use a GA to learn, Wei et al.(2013*a*, 2013*b*) introduce a method to classify the limit order market conditions and find that the expectation of the uninformed traders improves when they learn from the lagged fundamental value, historical prices and the bid-ask midpoint. In this paper, we use fundamental value, market price and the order book information, including the bid-ask spread, the change in bid/ask, and the change of depth to generate 18 CRs listed in Table 1 to describe market conditions. The CRs in Table 1 are grouped based on three aspects of market information. The first group, CR1 to CR6, describes the current price dynamics, the second group, CR7 to CR11, describes the recent change in the limit order book, and the third group, CR12 to CR18, describes the recent change in the depth of the limit order book[6]. We use binary strings to represent CRs and hence market condition. For example, "1" indicates that CR1 is true and "0" means that CR1 is false. Hence one binary string has 18 bits and every

[6]If we let the agent consider all the information of the limit order book and prices of the past τ periods, the agent may learn better, but it leave the set of classify rules too large and the learning more complicated. In our model, agents are not fully rational rather they have bounded rationality as real human beings, so that they can process part of information of the limit order book. In the classifier system, the average price and the order book depth reflect part of information of past τ periods.

bit represents two states of true or false of each CR, for example, "101110 01101 0110010" indicates one possible market condition. In some cases, some market information becomes irrelevant and in such case we use "#" to replace 1 or 0, indicating that the corresponding market information is not considered, for example, "##1#10 0#1#1 0#1#0#1" represents one possible market condition.

Mark Number	Classified rule	Description
		Current price dynamic
CR1	$p_{t'}^{m} > v_t^j$	The bid-ask midpoint is higher than the expected fundamental value
CR2	$p_{t'} > v_t^j$	The market price is higher than the expected fundamental value
CR3	$\overline{p}_{t,\tau} > v_t^j$	The average market price of last tau periods is higher than the expected fundament value
CR4	$p_{t'} > \overline{p}_{t,\tau}$	The current market price is higher than the average market price of last tau periods
CR5	$a_{t'} > v_t^j$	The current ask is higher than the average market price of last tau periods
CR6	$b_{t'} > v_t^j$	The current bid is higher than the average market price of last tau periods
		Recent changes of limit order book
CR7	$p_{t'}^{\pm}$	Last transaction sign (buy or sell)
CR8	$s_{t'} > s_{t'-1}$	The Current spread is bigger than the last spread
CR9	$p_{t'} > p_{t'-1}$	The current market price is high than the last market price
CR10	$a_{t'} > a_{t'-1}$	The current ask is higher than the last ask
CR11	$b_{t'} > b_{t'-1}$	The current bid is higher than the last bid
		Order book depth
CR12	$d_{t'}^{a} > d_{t'-1}^{a}$	The current depth at the ask is larger than the last depth at the ask
CR13	$d_{t'}^{aa} > d_{t'-1}^{aa}$	The current depth above the ask is larger than the last depth above the ask
CR14	$d_{t'}^{b} > d_{t'-1}^{b}$	The current depth at the bid is larger than the last depth at the ask
CR15	$d_{t'}^{bb} > d_{t'-1}^{bb}$	The current depth below the ask is larger than the last depth below the ask
CR16	$d_{t'}^{a} > d_{t'}^{b}$	The current depth at the ask is larger than the current depth at the bid
CR17	$d_{t'}^{aa} > d_{t'}^{bb}$	The current depth above the ask is larger than the current depth below the bid
CR18	$d_{t'}^{s} > d_{t'}^{b}$	The current depth of the sell side is larger than the current depth of the buy side

Table 1. The classified rules (CRs) based on price dynamics (from CR1 to CR6), change in the book (from CR7 to CR11), and order book depth (from CR12 to CR18). For informed traders, $v_t^j = v_t$ while for uninformed traders, $v_t^j = v_{t-\tau}$.

The second component of a trading strategy is the action with respect to buy/sell and order aggressiveness. In general, a trader can have many types of orders to choose from. Goettler, Parlour and Rajan (2005) classify orders into four types, including market order, aggressive limit order, limit order at the quote, and limit order away the quote. In Menkhoff et al. (2010) orders are classified into market orders, aggressive limit orders, and patient limit orders (limit orders at the quote or limit orders away from the quote). In this paper, we also classify orders into four types: a market order (MO), a limit order at the quote (LOA), an aggressive limit order (ALO), and an unaggressive limit order (ULO). To simplify the analysis, we define an aggressive limit order (ALO) to be a limit order above the bid or below the ask by one tick size, and an unaggressive limit order (ULO) to be a limit order below the bid or above the ask by one tick size. Therefore ALO narrows the bid-ask spread and improves the liquidity, while the LOA does not narrow the bid-ask spread but supplies immediate liquidity. We list all the actions (order types) in Table 2. Given the two sided of the book and the four types of orders, there are 8 order types in total. We use three binary bits to describe actions, for example "000" means a market buy MB order. For the informed traders, since they know the fundamental value, they can buy low and sell high by comparing the fundamental value to the bid and ask, but they can use GA to optimize their order aggressiveness. For the uninformed trader, they can use the GA to optimize both order type (buy or sell) and order aggressiveness. By putting the two components together, a trading strategy (x^i, y^i) means to take an action y^i (order type) under certain market conditions x^i. For example, one possible trading strategy i can be defined when x^i is given by "##1#10 0#1#1 0#1#0#1" and y^i is given by "000". Obviously, in some special limit order book scenarios, certain types of actions or orders are impossible or unused, which are listed in Table 3.

Action (buy)	Binary code	Description	Action(sell)	Binary code	Description
MB	000	Market buy	MS	111	Market sell
ALB	001	Aggressive limit buy	ALS	110	Aggressive limit sell
LBB	010	Limit buy at the bid	LSA	101	Limit sell at the ask
ULB	011	Unaggressive limit buy	ULS	100	Unaggressive limit sell

TABLE 2. The order types.

Scenario		Unused action
The limit order book is not empty	The bid-ask spread is more than one tick size	None
	The bid-ask spread is equal to one tick size	ALB & ALS
The limit order book is empty	Only when the buy limit order book is empty	MS, ALB & ULB
	Only when the sell limit order book is empty	MB, ALS & ULS
	Both the buy and the sell limit order books are empty	MB, ALB, ULB, MS, ALS & ULS

TABLE 3. Limit order book scenarios and the unused actions.

We now turn to the evolution process of the GA, including selection, crossover and mutation. In the selection process, a chromosome is selected by a tournament mechanism based on its strength S^i. Initially, all chromosomes are randomly generated and the market conditions x^i and strengths S^i of all chromosomes are the same. Traders randomly select chromosomes in the early simulation periods. When a submitted order of a trader has been executed, or canceled, or expired, the trader updates the performance and then the strength of the chromosome. The strength $S^i_{t'}$ at time t' is equal to the performance $\pi^i_{t'}$ minus the specificity $\delta^i_{t'}$ measuring the cost of the chromosome.[7] We use the order profit $r^i_{t'}$ to measure the performance of the chromosomes i. Hence $r^i_{t'} = v_{t'} - p_{t'}$ for an executed buy order and $r^i_{t'} = p_{t'} - v_{t'}$ for an executed sell order. If the order has been canceled or it has expired, $r^i_{t'} = 0$. For the informed traders, their performances are updated immediately when they enter the market. For the uninformed traders, due to the information lag, their performances are updated only when the transactions occur before or at period $t - \tau$. The performance of the chromosome i is updated according to a weighted average of the recent performance and historical performance

$$\pi^i_{t'} = \beta r^i_{t'} + (1 - \beta)\pi^i_{t'-1}, \tag{1}$$

where $\beta \in [0, 1]$ is the weight of the recent performance. A larger β means that traders put more weight on the recent performance and less weight on the historical performance.

After the selection process of a trader, new chromosomes are generated through the processes of crossover and mutation according to given probabilities. Crossover means that, with a certain probability called the crossover rate, the trader randomly chooses two high strength chromosomes as parents, splits each chromosome into two parts at a random bit and then swaps the two parts to create two new chromosomes as children. This process is illustrated in Figure 1. The two parents chromosome are "10011" and "01100". If they are split at the third bit, then two new child chromosomes are "01111" and "10000". The strength of the child chromosomes are equal to the average strength of their parents. Mutation means that, with a certain probability called the mutation rate, the trader randomly selects a high strength chromosome as a parent and makes a random bit change of the parent chromosome to a different value. This is also illustrated in Figure 1. For the parent chromosome "10011", the second bit is chosen to mutate, then the child chromosomes become either "11011" or "1#011". The strength of the child chromosomes is equal the the strength of the parent chromosome minus its specificity.

In our simulations, the evolution process of a GA is active every 120 periods[8] When the evolution process is active, traders delete a proportion[9] of low strength

[7]The specificity measures the fitness or cost of the chromosome. For example, for a chromosome "1#0#1", the number of specific bits (non-omitted bits) m is equal to 3. The specificity of a chromosome is equal to $m\mu$, where μ is the bit cost and is set to be a small value, such as 0.001. Hence, $S^i_{t'} = \pi^i_{t'} - \delta^i_{t'}$, meaning that, if two chromosomes have the same performance $\pi^i_{t'}$, the chromosome with less specificity (with more omitted bits) has a higher strength.

[8]If a simulation runs for 36,000 periods, the evolution process has been active for 300 times.

[9]We set the proportion to be 10%.

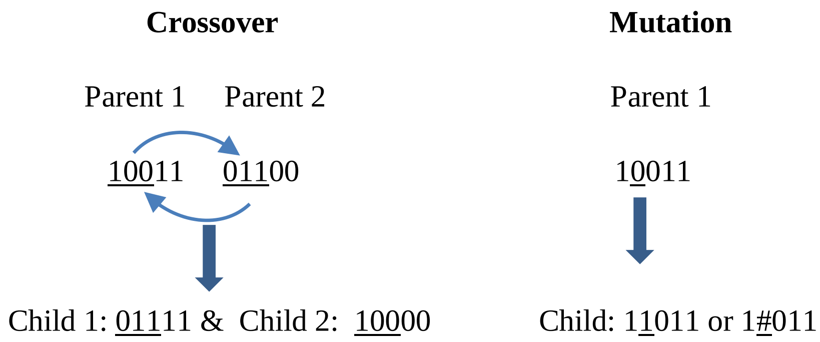

FIGURE 1. The crossover and mutation processes of a genetic algorithm.

chromosomes and replace them by new chromosomes which have been created by the evolution process.

2.3. **Experimental design, parameter setting, and performance measures.** We take the two-sided learning model as a benchmark model (BM) in which both informed and uninformed traders use GA learning. In order to examine the effect of learning, we consider a case (A1) of no learning for all traders and two cases of one-sided learning in which only the informed (A2) or the uninformed (A3) traders learn. These three scenarios (A1, A2 and A3) are grouped into group A. By comparing A2 (A3) to the BM, we can examine the learning effect of the uninformed (informed) traders from an one-sided learning mechanism to a two-sided learning mechanism. Intuitively, the informativeness of the market, the volatility of the fundamental values, the information-lived time, and the discount rate of the historical performance in the GA may have a significant impact on the market. Therefore, as a robustness test, we design several experiments to examine these effects. In particular, groups B, C, D and E aim to examine the effect of the informative level (group B), the volatility of fundamental value (group C), the information-lived time (group D), and the historical performance evaluation (group E), respectively. All the experiments are listed in Table 4.

For the parameters, we choose the initial fundamental value $v_0 = 20$ and the initial market price $p_0 = v_o = 20$. The tick size is 0.02, corresponding to 0.1% of the initial market price. The total population of the market is set to 1000. Based

Experiment	Informed	Uninformed	κ	τ	β
BM	100	900	5	60	0.6
A1	100 ***random***	900 ***random***	5	60	0.6
A2	100	900 ***random***	5	60	0.6
A3	100 ***random***	900	5	60	0.6
B1	***1***	***999***	5	60	0.6
B2	***500***	***500***	5	60	0.6
C1	100	900	***2***	60	0.6
C2	100	900	***10***	60	0.6
D1	100	900	5	***30***	0.6
D2	100	900	5	***120***	0.6
E1	100	900	5	60	***0.1***
E2	100	900	5	60	***0.9***

TABLE 4. The experiment design and parameters. In A1,A2 and A3, "random" means that this type trader does not employ GA learning and randomly chooses his/her actions.

on some empirical studies on the probability of informed trading (PIN),[10] we set the proportion of the informed traders to be 10% in the BM, which corresponds to 100 informed traders and 900 uninformed traders. The volatility of the fundamental value in the BM is set as in Hollifield, Miller, Sandas and Slive (2006) and Goettler et al. (2009).[11] In the BM, we choose the Poison rate $\phi = 0.5$ and $\kappa = 5$. This implies that, on average, the innovation of the fundamental value occurs once every two minutes and each innovation changes the fundamental value by five tick sizes[12]. For the information-lived time, we set $\tau = 60$ in the BM. Considering that one period corresponds to one minute, $\tau = 60$ means that the time-lag of the fundamental values is one hour for the uninformed traders. We set the maximum order survival time as a trading day, hence $T = 240$. For the traders arrival rate, we assume that the informed and uninformed traders follow the same Poison process with arrival rate of

[10]The PIN models in the literature, such as Easley, Hvidkjaer and O'Hara (2010), Lin and Ke (2011) and Yan and Zhang (2012), find that the PIN is between 10% to 20% on average. Since the informed trader may have more transactions than uninformed traders, the proportion of informed trader may be less than the PIN level. Therefore we set the proportion of the informed traders to be 10%.

[11]Hollifield et al. (2006) find that the expected variation of the fundamental value is 1.7% over a 10-minute interval on the Vancouver exchange. Goettler et al. (2009) follow their empirical findings.

[12]Following this innovation process,the fundamental value may be negative at sometimes. So we set the minimum fundamental value to 5 to avoid a negative state.

$\lambda = 0.017$, meaning that each trader enters the market four times a day on average.[13] For the discount rate of historical performance, we set $\beta = 0.6$ in the BM. Also the crossover rate or probability is 0.1 and the mutation rate or probability is 0.3. The parameter settings for other experiments are the same as in the BM, expect for the specifical values listed in Table 4.

To measure the impact of learning on price dynamics, we employ the convergence of the market price to the fundamental value to measure the informational efficiency. Following Theissen (2000), we use Mean Absolute Error (MAE) to measure the deviation or error of the market price p_t from the fundamental value v_t,

$$MAE = \frac{1}{T}\sum_{t=1}^{T} |p_t - v_t|. \tag{2}$$

We also use Mean Relative Error (MRE) to measure the relative error of the market price from the fundamental value,

$$MRE = \frac{1}{T}\sum_{t=1}^{T} \frac{|p_t - v_t|}{v_t}. \tag{3}$$

In order to obtain statistical significance we run 30 simulations. Since the GA needs sufficient learning time to obtain optimal trading strategies, each simulation runs 48,000 periods, but the analysis is based on the results from periods 36,001 to 48,000, in total of 12,000 periods, which is about 200 hours.

3. Price dynamics

This section examines the impact of learning on the price dynamics. The analysis is focuses on informational efficiency and some stylized facts in limit order markets.

3.1. **Price convergence and informational efficiency.** Intuitively, when learning is effective, it helps market informational efficiency and the convergence of the market price to the fundamental price. This is supported by typical simulations in Figure 2, in which both market price and the fundamental price over the first 4500 periods are plotted for the BM and Group A experiments. It illustrates that, with no learning (A1), the market prices deviate away from the fundamental values. When only the informed traders learn, A2 shows that the market prices do not converge to the fundamental values but the deviations are much less than in the case A1. When only the uninformed traders learn, A3 shows that market prices gradually converge to the fundamental value. Finally, in the BM case of two-sided learning, the market prices converge to the fundamental value very quickly.

[13]We assume that traders do not continuously monitor the market,Dugast (2012) argues that agents have limited attention and thus do not monitor the market continuously.

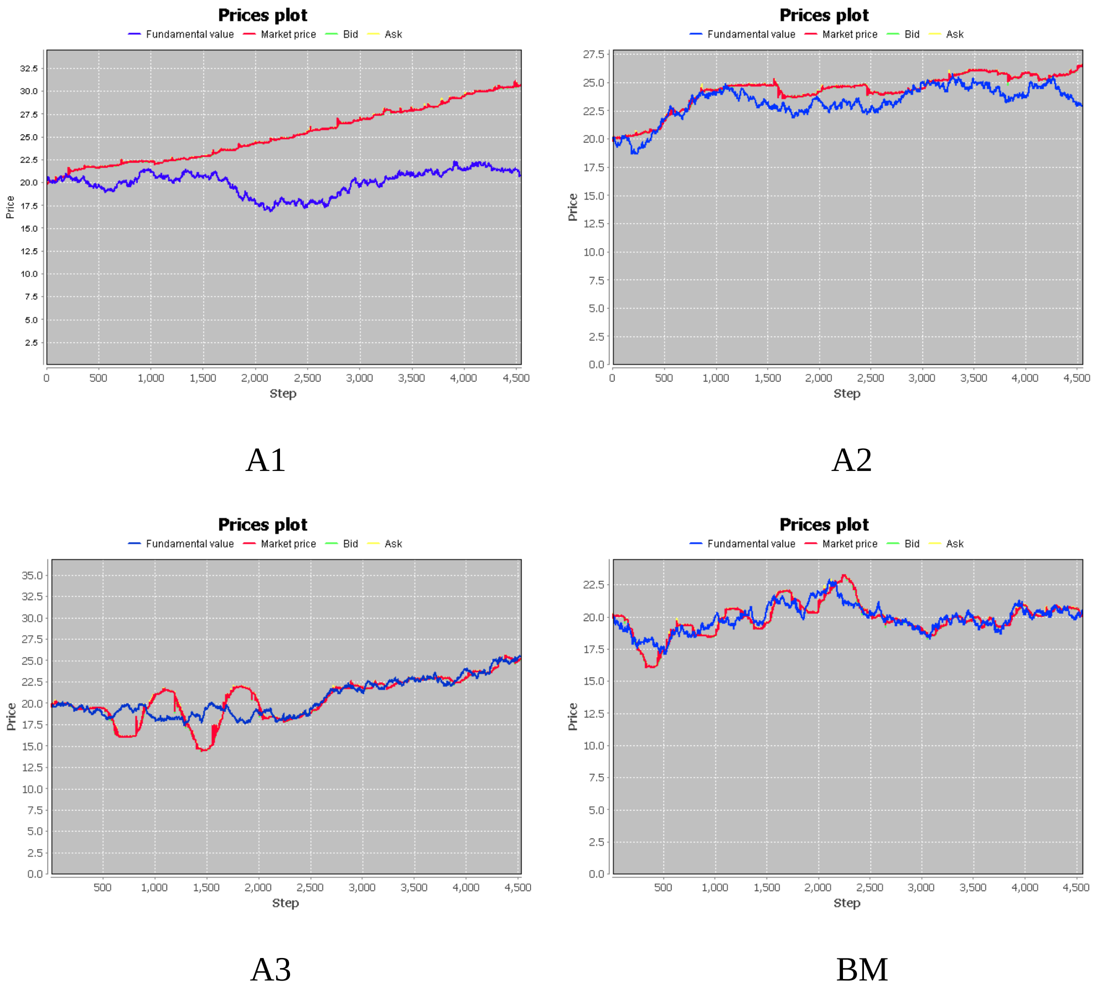

FIGURE 2. Price convergence in Group A and the experiment BM. The price plots are randomly chosen from 30 simulation runs in each experiment.

We further analyze the informational efficiency based on 30 simulations for each experiment and the results are summarized in Table 5. It shows that the informational efficiency is consistent with the price convergence illustrated in Figure 2. For A1, both MAE and MRE are very large. Both MAE and MRE are reduced for A2 and significantly reduced for A3. For BM, both MAE and MRE are the smallest among all cases. This implies that learning of both the informed and uninformed traders contribute to price convergence and informational efficiency[14]. However the learning from the uninformed traders is more effective that from the

[14]It means that the price often converges to the fundamental value, but not always converges to the fundamental value all the time.

informed traders. Table 5 also shows that informational efficiency is positively related to the proportion of the informed traders (Group B) but negatively to the volatility of fundamental value (Group C) and the information-lived time (Group D). These results are consistent with the one-sided learning model in Wei, Zhang, He and Zhang (2013*b*). The impact of the performance weight β on the informational efficiency is more complicated. Comparing the results in Group E to the BM, it seems to indicate that extreme values in β is not optimal for informational efficiency.

Experiment	*MAE*	*MRE*
BM	0.42	2.29%
A1	122.28	912.12%
A2	38.26	194.30%
A3	0.43	3.02%
B1	0.75	5.12%
B2	0.13	0.72%
C1	0.15	0.82%
C2	1.00	4.52%
D1	0.37	2.26%
D2	0.48	2.70%
E1	0.42	2.42%
E2	0.48	2.47%

TABLE 5. Informational efficiency indicators based on the average of 30 simulations.

3.2. **Stylized facts.** The stylized facts of empirical studies of limit order markets include leptokurtosis, fat tails, short-term autocorrelation, volatility clustering, and long memory with positive long-range correlation (see Gould, Porter, Williams, Fenn and Howison (2012). Wei, Zhang, He and Zhang (2013*a*) show that the one-sided learning model can generate all these stylized facts, among which long memory is the most important one. Hence we focus on the long memory by examining the Hurst exponent H. According to Gould et al. (2012), return series have long memory with positive (negative) long-range autocorrelations when the Hurst exponent satisfies $0.5 < H < 1$ ($0 < H < 0.5$). If $H = 0.5$, then return series follow a random walk. Table 6 shows that the return of the market price can generate long memory with positive long-range correlation when uninformed traders learn. This result is consistent with Wei et al. (2013*a*). This stylized fact also reflects the fact that both the market price and mid-price are not martingale processes and the prices are not efficient prices as in Goettler et al. (2009) and Rosu (2012*a*), but their dynamics closer to the real markets.

Experiment	H_{p_t}	$H_{p_t^m}$	H_{v_t}
BM	0.56	0.73	0.53
A1	0.20	0.48	0.52
A2	0.34	0.56	0.53
A3	0.61	0.76	0.53
B1	0.62	0.77	0.54
B2	0.43	0.56	0.53
C1	0.52	0.64	0.54
C2	0.59	0.80	0.54
D1	0.58	0.74	0.53
D2	0.58	0.74	0.53
E1	0.59	0.75	0.52
E2	0.58	0.75	0.54

TABLE 6. The Hurst exponents H_{p_t} and $H_{p_t^m}$ based on market price, mid-price and H_{v_t} based on the fundamental values.

In summary, learning from traders, in particular uninformed traders, can improve market informational efficiency and generate realistic stylized facts in the limit order market.

4. Limit order book phenomena

This section focuses on some limit order book phenomena, such as order submission, bid-ask spread, order book shape and order profit. They reflect market liquidity and whether traders consume or supply the liquidity.

4.1. **Order submission.** For order submission, we mainly focus on the order aggressiveness, liquidity consumption and supply. Since our model allows traders to endogenously choose their order aggressiveness, we can classify orders from type j traders ($j = I$ for the informed traders and $j = U$ for the uninformed traders) to market order MO_j, unaggressive limit order ULO_j, limit order at quote LOA_j, and aggressive limit order ALO_j, where

$$MO_j = MB_j + MS_j, \qquad ULO_j = ULB_j + ULS_j,$$
$$LOA_j = LBB_j + LSA_j, \qquad ALO_j = ALB_j + ALS_j.$$

We reported the total order numbers in Table 7, together with the executed limit orders ELO_j for the informed (the left panel) and uninformed (the right panel) traders.

Experiment	ULO_I	LOA_I	ALO_I	MO_I	ELO_I	ULO_U	LOA_U	ALO_U	MO_U	ELO_U
BM	5,853	5,832	3,035	5,744	4,719	50,884	43,689	34,940	54,068	56,210
A1	5,842	5,846	2,955	5,809	5,162	52,361	39,518	39,460	52,397	57,221
A2	4,784	4,815	3,032	7,714	5,039	52,982	38,857	38,760	53,029	60,014
A3	5,846	5,836	2,936	5,767	4,657	50,972	43,723	34,702	54,212	56,414
B1	58	57	29	58	51	56,710	48,323	37,901	60,872	62,212
B2	28,095	28,024	18,434	27,416	21,232	27,245	24,564	21,963	28,300	34,731
C1	5,934	5,892	2,747	5,822	4,867	52,033	43,614	33,297	54,564	56,733
C2	5,738	5,741	3,387	5,567	4,537	49,507	43,854	36,774	53,490	55,415
D1	5,823	5,805	3,091	5,677	4,692	50,513	43,535	35,270	54,287	56,357
D2	5,827	5,820	2,997	5,722	4,740	51,113	43,913	34,487	54,033	56,120
E1	5,844	5,839	2,977	5,760	4,660	51,600	45,069	33,521	53,392	55,400
E2	5,886	5,816	3,009	5,758	4,764	51,266	43,698	34,560	54,089	56,198

TABLE 7. Total numbers of different types of orders for the informed and uninformed traders.

To measure the order aggressiveness and liquidity supply and demand, we introduce the following ratios:

$$AL_j = \frac{ALO_j}{LO_j}, \qquad AGG_j = \frac{ALO_j + MO_j}{LO_j + MO_j},$$

$$SUB_j = \frac{LO_j}{LO_j + MO_j}, \qquad TAK_j = \frac{MO_j}{MO_j + ELO_j}, \qquad LE_j = \frac{ELO_j}{LO_j},$$

where

$$LO_j = ULO_j + LOA_j + ALO_j$$

is the total number of limit orders and $j = I, U$. Therefore, for type j traders, AL_j measures the ratio of aggressive limit orders to total limit orders and AGG_j measures the order aggressiveness. Also, SUB_j measures the order submission rate, TAK_j measures the taking rate, and LE_j measures the limit order execution rate. Following Bloomfield et al. (2005), we use the order submission rate SUB_j to measure the liquidity supply and the taking rate TAK_j to measure liquidity consumption. Based on 30 simulations for each experiment, the results are reported in Table 8.

To see the effect of learning, we consider first one-sided learning and obtain the following observation.

Observation 1. *The one-sided learning from the informed traders increases their order aggressiveness and makes them consume the liquidity, while the effect of one-sided learning from the uninformed traders becomes less significant (although they submit less aggressive limit orders and consume the liquidity).*

When the informed traders learn but not the uninformed traders, we compare A1 and A2 in Table 8 and find that AL_I increases from 20.18% to 24.00%, AGG_L increases from 42.85% to 52.82%; TAK_I increases from 52.95% to 60.49%, while SUB_I decreases from 71.60% to 62.08%. These results imply that the informed traders become more aggressive, in particular they submit more market orders to

Experiment	AL_I	AGG_I	SUB_I	TAK_I	LE_I	AL_U	AGG_U	SUB_U	TAK_U	LE_U
BM	20.62%	42.90%	71.93%	54.90%	32.06%	26.98%	48.48%	70.55%	49.03%	43.40%
A1	20.18%	42.85%	71.60%	52.95%	35.25%	30.04%	49.99%	71.48%	47.80%	43.57%
A2	24.00%	52.82%	62.08%	60.49%	39.89%	29.68%	49.99%	71.12%	46.91%	45.95%
A3	20.08%	42.69%	71.71%	55.32%	31.86%	26.82%	48.43%	70.47%	49.00%	43.60%
B1	20.14%	43.07%	71.29%	53.21%	35.42%	26.52%	48.46%	70.13%	49.46%	43.52%
B2	24.73%	44.96%	73.11%	56.36%	28.48%	29.77%	49.24%	72.27%	44.90%	47.08%
C1	18.85%	42.02%	71.45%	54.47%	33.40%	25.82%	47.88%	70.27%	49.03%	44.00%
C2	22.78%	43.82%	72.75%	55.10%	30.52%	28.26%	49.16%	70.87%	49.12%	42.58%
D1	21.00%	42.99%	72.17%	54.75%	31.88%	27.27%	48.78%	70.43%	49.06%	43.58%
D2	20.47%	42.81%	71.90%	54.69%	32.37%	26.63%	48.23%	70.56%	49.05%	43.33%
E1	20.31%	42.79%	71.79%	55.28%	31.79%	25.75%	47.34%	70.92%	49.08%	42.55%
E2	20.45%	42.83%	71.87%	54.72%	32.38%	26.68%	48.28%	70.54%	49.04%	43.39%

TABLE 8. Ratios of order submission.

consume liquidity. However, when only the uninformed traders learn, we compare A1 and A3 and find that AL_U decreases significantly from 30.04% to 26.82%, also AGG_U decreases from 49.99% to 48.43%, meaning that the uninformed traders reduce their order aggressiveness. TAK increases from 47.80% to 49.00% (meaning they consume liquidity), while SUB_U deceases from 71.48% to 70.47%, however, the changes are less significant.

We now explore the effect of two-sided learning and obtain the following observations.

Observation 2. *Comparing one-sided learning to two-sided learning, the impact of the learning from the informed traders on order submission for both types of traders is not significant, but becomes significant when the uninformed traders learn. Also learning from the uninformed traders can make the informed traders reduce their order aggressiveness and supply more liquidity, but make themselves consume the liquidity by submitting more market orders or limit orders at the quote.*

Indeed, when the informed trader learn, we compare A3 and BM and find that all the ratios change slightly. However, when the uninformed traders learn, we compare A2 and BM. On the one hand, for the informed traders, AL_I decreases from 24.00% to 20.62%, and AGG_I decreases significantly from 52.82% to 42.90%; TAK_I decreases from 60.49% to 54.90% while SUB_I increases from 62.08% to 71.93%. Hence the learning from the uninformed traders make the informed traders reduce order aggressiveness and supply more liquidity. On the other hand, for the uninformed traders themselves, AL_U decreases from 29.68% to 26.98%. Both AGG_U and SUB_U decrease slightly, but TAK_U increases from 46.91% to 49.03%. If we check the total number of orders in Table 7, we can see that ALO_U decrease from 38,760 to 34,960, but MO_U increases from 53,029 to 54,068. This implies that the uninformed traders use market orders to replace some aggressive limit orders, thus they consume more liquidity. In addition, ULO_U decreases but LOA_U increases significantly from 38,857 to 43,689. This indicates that the uninformed traders use more limit orders at quote to replace some aggressive limit orders. This finding is a supplement to Menkhoff et al. (2010) who find that uninformed traders only use aggressive limit orders to replace patient limit orders (which equals to ULO_U plus LOA_U).

Observation 3. *When both types of traders learn, the informed traders reduce aggressive limit orders but increase market orders than the uninformed traders so that the informed traders have more willingness to consume liquidity than the uninformed traders, although their liquidity supply is about the same.*

In fact, in all the two-sided learning cases in Table 8, AL_I is less than AL_U, but AGG_I is higher than AGG_U. This indicates that, although the informed traders submit less aggressive limit orders, but they submit more market orders, becoming more aggressive than the uninformed traders and more prefer to consume liquidity. This result is intuitive because when information is short-lived, the informed traders need to be more aggressive and trade more in order to benefit from their information advantage. Our result is different from observations in Goettler et al. (2009) and Rosu (2012*a*) for different reasons. For Goettler et al. (2009), the private value is the most important factor for order submission, the uninformed traders with high private value prefer to submit market orders, so that when informed traders set limit orders at more profitable prices, limit orders still have high executed probability, thus informed traders with zero private value prefer to use limit orders and supply liquidity; in our model, there is no private value, uninformed traders submit orders based on the market conditions, so that the limit orders are not so profitable for the informed traders, and informed traders face pick-off risk due to the competition with other informed traders and the learning of uninformed traders, thus when they learn, they find that using more market orders can obtain more order profit. However, when the volatility of fundamental value becomes higher, informed traders' order submission is similar to our model (see Observation 5 later). For Rosu (2012*a*), when the information is long-lived, informed traders are patient so that they obviously prefer to use limit orders to supply liquidity.

Observation 4. *When market becomes more informative, both the informed and uninformed traders increase their order aggressiveness and liquidity supply; however the informed traders increase liquidity consumption while the uninformed traders reduce consumption of liquidity.*

This observation is based on the comparison between B1 and B2. It is found that, as the number of the informed traders increases, AL_L, AL_U, AGG_I, AGG_U, SUB_I, and SUB_U all increase, and TAK_I increases but TAK_U decreases. This result shows that when the market becomes more informative, the uninformed traders prefer to increase liquidity supply and reduce liquidity consumption, but the informed traders increase their liquidity supply and consumption, so that they become more active because of the high level of competition among themselves.

Observation 5. *When the volatility of fundamental value increases, the informed traders increase order aggressiveness and consumption.*

This is consistent with Goettler et al. (2009), when the volatility of the fundamental value becomes higher, informed traders prefer to use market orders to reduce pick-off risk and make more profit by finding mispriced orders in the limit order book. But in our model, informed traders also submit more aggressive limit orders to supply liquidity. The uninformed traders also increase their order aggressiveness, but the impact on the liquidity is less significant (although their liquidity supply and consumption increase slightly). This observation is based on the comparison between C1 and C2. This result is consistent with Menkhoff et al. (2010) in that the uninformed traders are relatively insensitive to the volatility.

Observation 6. *When the information lag increases, both the informed and uninformed traders reduce their aggressive limit orders slightly; but the order aggressiveness, liquidity consumption and supply do not change much.*

This is based on a comparison between D1 and D2. Therefore, the information-lived time does not have a significant impact on the market liquidity.

Observation 7. *When the uninformed traders put more weight on the recent performance, there is no significant impact on liquidity, although the uninformed traders increase order aggressiveness slightly by using more aggressive orders.*

Overall, we can see that learning from the traders, in particular from the uninformed traders, can affect the order submission behavior and consumption and supply of market liquidity of both types of traders significantly.

4.2. **Order book shape.** Following the previous discussion on order submission behavior, we now examine the shape of the order book. We have obtained the following observations.

Observation 8. *The order book displays a "hump" shape near the quote and the depth of the order book is peaked at the second best quote on both side of the book.*

This is illustrated in Figure 3 for BM and Group A experiments. Actually, the hump shape is observed across all experiments, reflecting the fact that the hump shape order book near the quote is generated by the continuous double auction trading mechanism and the order types. This is consistent with the literature and empirical observation. Figure 4 displays the order book shape based on one month high frequency data of the Westpac Bank Corporation (WBC), a bank stock in the Australian stock market in June 2012. For the relative limit order price levels 1 to

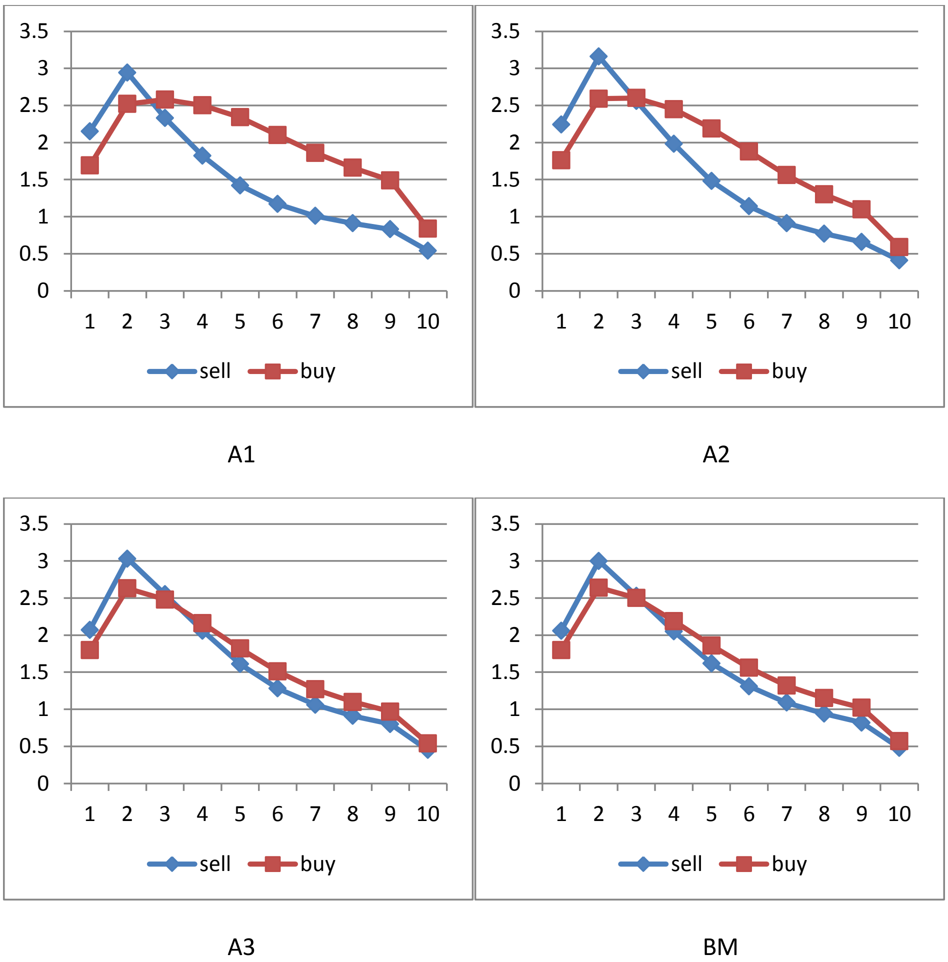

FIGURE 3. Order book shape and the effect on learning.

10, the order book shape in Figure 3 is similar to the one in Figure 4. However, the model is not able to generate the hump shape far away from the quote. That is partially due to the simplicity of the unaggressive limit orders with only one tick size away the quote in our experiments, partially due to the limitation of trading only one share each time and order cancelation on re-entering. In real markets, traders can submit unaggressive limit orders with many relative price levels due to their private preferences, such as patient or impatient and optimism or pessimism.[15]

Observation 9. *Learning of the traders, in particular the uninformed traders, reduces order imbalance beyond the quote.*

This observation is based on a comparison among A1 (no learning), A2 (only the informed traders learn), A3 (only the uninformed traders learn), and BM (both

[15]Rosu (2009) conjectures that patient traders optimistically consider limit orders that are far away from the bid and ask which might be executed in the future, while impatient traders pessimistically believe that these order will never be matched and executed.

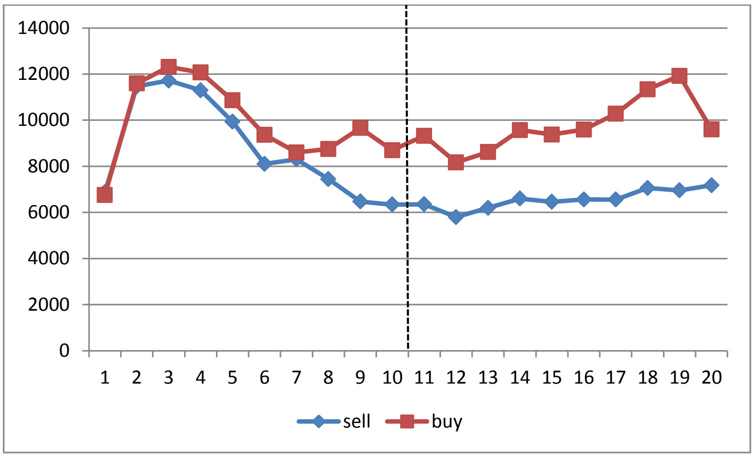

FIGURE 4. The limit order book shape of the Westpac Bank Corporation in the Australian Stock Market based on one month tick by tick data in June, 2012 from SIRCA.

the informed and uninformed traders learn) and the results are reported in Figure 3. We can see that the order book shape changes significantly from A1 to A3 but insignificantly from A1 to A2 and from A3 to BM. This implies that the impact of learning on the order book shape is stronger for the uninformed than the informed traders. In addition, comparing A3 to BM, the learning from the informed traders has little impact on the order book shape. This observation is consistent with the previous order submission analysis on the significant impact of learning from the uninformed traders. It can be explained by the evolution of trading strategies (see Observation 16 later).

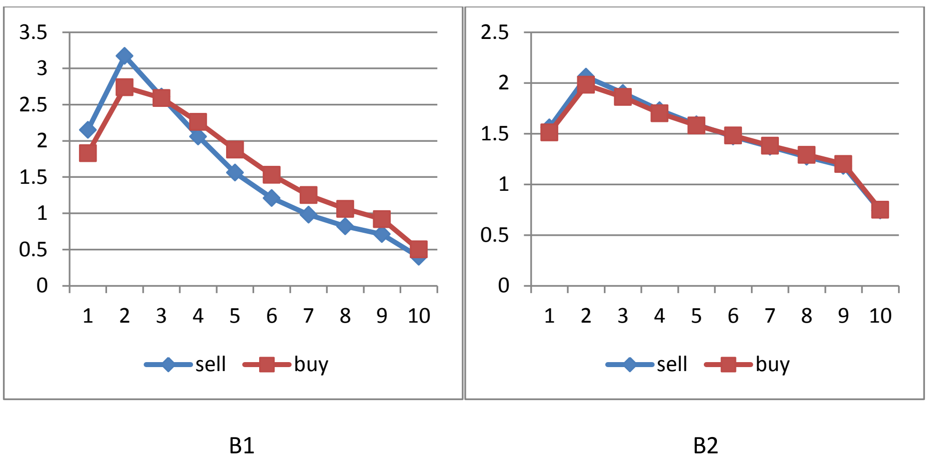

FIGURE 5. The impact of the proportion of informed traders on the limit book shape.

Observation 10. *The order book becomes thinner when the market becomes more informative.*

This observation is based on a comparison between B1 and B2 in Figure 5. From B1 to B2, the depth of order book across quotes is reduced uniformly. Intuitively, when there are more informed traders, they tend to submit more market orders and consume liquidity, which then reduces orders in the book.

Further analysis finds that the volatility of the fundamental value, the information-lived time, and the performance weight have an insignificant impact on limit order book shape (see Table 12 in Appendix A).

4.3. **Bid-ask spread.** The bid-ask spread is showed in Table 9. We have two findings.

Observation 11. *The one-sided learning from either the informed or uninformed traders narrows the bid-ask spread and reduces the spread volatility. For two-sided learning, learning from the uninformed traders also narrows the spread and reduces the volatility, but the impact is opposite under the learning of the informed traders.*

This observation is based on comparisons among no learning (A1), one-sided learning (A2 and A3) and two-sided learning (BM) and the results in Table 9. When the uninformed traders learn, both the spread and volatility from A1 to A3 and from A2 to BM are reduced. This is also the case for one-sided learning from the informed traders by comparing A1 to A2, but is opposite for two-sided learning by comparing A3 and BM. The main reason is that the learning from the uninformed traders make the market price much closer to the fundamental value and hence the impact on the spread is mainly due to the impact on price rather than the order submission (see the change of ALO_j and MO_j in Table 7).

Experiment	Spread	Variance
BM	4.4	7.2
A1	16.4	33.8
A2	6.4	13.8
A3	3.6	5.7
B1	3.4	5.3
B2	6.4	8.5
C1	2.4	3.4
C2	8.1	13.4
D1	4.1	6.4
D2	4.1	6.6
E1	3.7	5.8
E2	3.9	6.3

TABLE 9. Bid-ask spread and its volatility in tick sizes.

Observation 12. *When both the informed and uninformed traders learn, the increase in the proportion of informed traders, the volatility of fundamental value, the information-lived time, and the performance parameter tend to increase the bid-ask spread and spread volatility.*

This observation is based on the results in Table 9 for Groups B, C, D and E. The result is consistent with Wei et al.(2013*a*, 2013*b*). It is found that the order submission can explain the bid-ask spread changes in D and E, but not in Group C[16]. By comparing D1 and D2, E1 and E2, we can see that the change in the aggressive limit orders and the market orders are insignificant, which explains the insignificant change in the bid-ask spread. However, if we compare C1 and C2, the change of the aggressive limit order and the market orders can not explain the change in the bid-ask spread. In this case, the aggressive limit order increases by 4,117 shares and the market orders decreases by 1,329 shares. Intuitively, the bid-ask spread should be narrowed, however the bid-ask spread actually increases. This is due to the effect of price convergence. As we have shown in Table 5, both MAE and MRE for C2 are much larger than for C1. This implies that the market price deviates from the fundamental value significantly, which then increases the bid-ask spreads.

In summary, it is found that, when the price deviation is large, the bid-ask spread is determined by the learning ability of the traders for price convergence. However when the price deviations are small, the bid-ask spread is determined by the order aggressiveness of the traders.

4.4. **Order profit.** Since it is a zero-sum game, the informed traders are expected to gain from the uninformed traders, however, effective learning from the uninformed traders should help them to reduce their loss significantly. This intuition is confirmed by the results on order profits in Table 10. It also shows that, when the proportion of the informed traders increases, they make less order profit due to the competition among themselves. When the volatility of fundamental value increases, the informed traders make more order profit due to their information advantage. However, we do not observe a significant effect of the information-lived time and performance weight parameter.

5. The evolution of trading strategies

To provide more insight into how the traders learn, we examine the evolution process of the trading strategies of the traders, in particular which classified rules

[16]For Group B, B2 is a special case in which the bid-ask spread is highly influenced by the exogenous setting of the bid and ask when the order book is empty. We assume that when the one-sided order book is empty, the bid-ask spread is about 1% of the market price. In B2, the order book is very thin (see Figure 5), sometimes the limit order book is empty, so that makes the bid-ask spread become much wider.

Experiment	r_I	r_U
BM	0.3927	-0.0374
A1	122.2717	-12.2375
A2	38.3874	-4.3348
A3	0.4069	-0.0385
B1	0.7300	-0.0008
B2	0.0452	-0.0349
C1	0.1369	-0.0133
C2	0.9307	-0.0864
D1	0.3345	-0.0315
D2	0.4538	-0.0432
E1	0.3879	-0.0372
E2	0.4508	-0.0431

TABLE 10. Order profit per order of the informed (r_I) and uninformed (r_U) traders.

have been often used when traders submit their orders. We record the relative frequency of the classified rules (CR) in the trading strategies, which have been selected when making trading decisions. We report the average frequency of CRs per trader and present the results in Figures 6 to 11 across all the experiments. The analysis helps us to understand how the traders process private and market information and how they interact with each other.

5.1. **The evolution of trading strategies for the informed traders.** First we examine the evolution process for the informed traders. Figures 6 to 8 show that all the classified rules are useful for the informed traders, however some rules are used more often than others, depending on market environment.

Observation 13. *The learning from the uninformed traders increases the use of all the CRs for the informed traders, in particular the CRs related to the price trend and the sign of the last transaction.*

This observation is based on a comparison between A2 and BM and the plot in Figure 6. It shows that the using frequency of the informed traders per trade increases significantly (by about 25%) across all the CRs, in particular, CR4 (the current market price is higher than historical price) and CR7 (the sign of the last transaction). This is because the learning of the uninformed traders highly affects the market price. Hence the informed traders can use the market price information to optimize their order aggressiveness and improve their profit opportunity.

Observation 14. *When the market becomes less informative, the informed traders care more about the imbalance between the current depth at the ask and the current depth at the bid but less about other depth changes.*

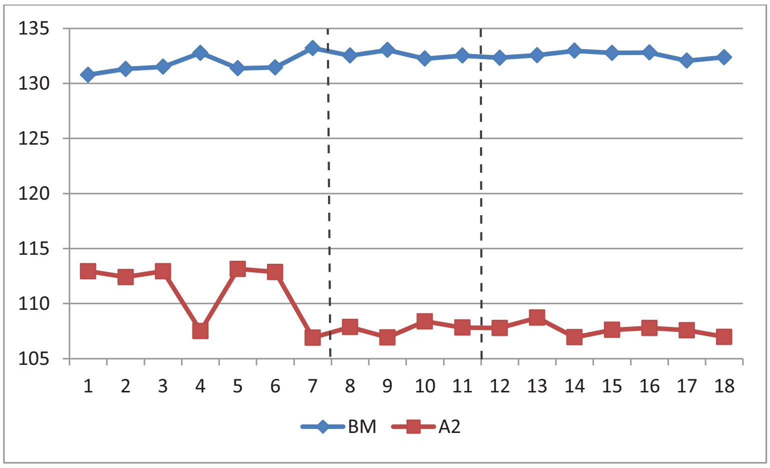

FIGURE 6. The average frequency of the classified rules CR1 to CR18 used by the informed traders per trader in experiments A2 and BM.

This observation is based on the two extreme market structures in Group B presented in Figure 7. It shows that the use of the CRs in B2 are similar. However CR16 is used more while CR12, CR15, CR17 and CR18 are used less in B1. In this case, there is only one informed trader. Due to the lack of competition, the informed trader cares more about the imbalance between depth at the bid and depth at the ask so that he can better forecast order execution.

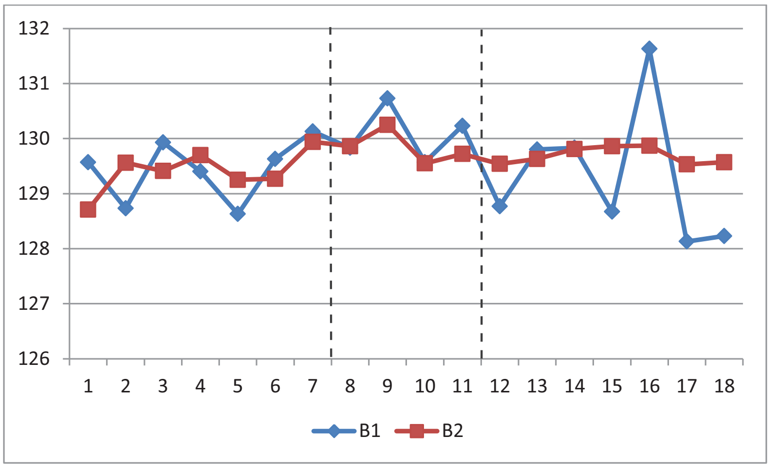

FIGURE 7. The average frequency of the classified rules CR1 to CR18 used by the informed traders per trade in experiments B1 and B2.

Observation 15. *In general, the informed traders care more about the market price change and the change of depth at the quote.*

Figure 8 presents the use of frequencies for BM and Groups C, D and E. They seem to share a similar pattern. Note that the use of CR4, CR7, CR9, CR14 and CR16 is higher than other CRs. CR4, CR7 and CR9 describe the changes of market price, while CR14 and CR16 describe the depth at the quote. This result is robust with respect to the change of the fundamental value volatility, the information-lived time, and the performance parameter. However, when traders put more weight on the recent performance, Figure 8 shows that the use of all the CRs are the highest. Intuitively, when traders care more about the recent performance, they submit more aggressive orders. Hence they care more about the current limit order book information.

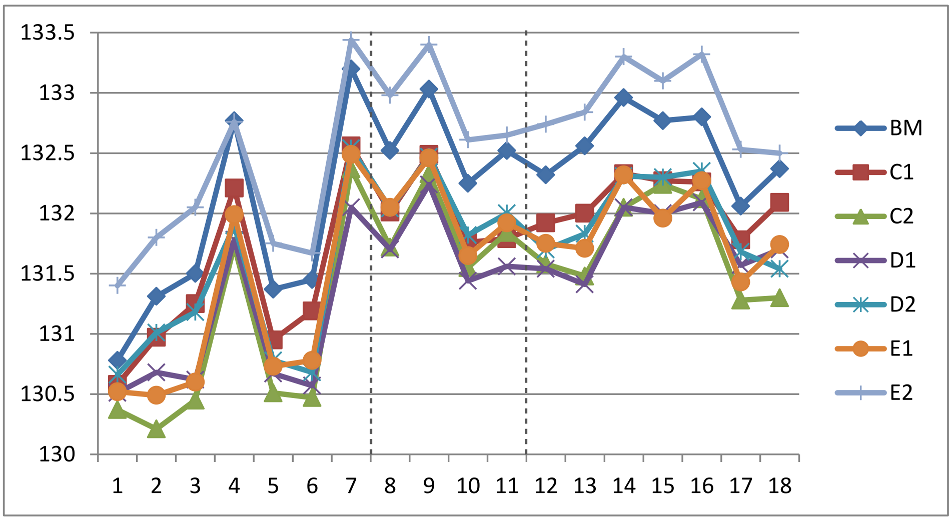

FIGURE 8. The average frequency of the classified rules CR1 to CR18 used by the informed traders per trade in experiments BM, Group C, D and E.

5.2. **The evolution of trading strategies for the uninformed traders.** Finally, we examine the evolution process of the uninformed traders.

Observation 16. *In general, the use of the CRs for the uninformed traders is less than for the informed traders. The learning from the informed traders does not affect significantly the use of the CRs for the uninformed traders, although they care more about the relation between the bid-ask midpoint and the fundamental value, and the imbalance of order book depth between buy and sell, but less about the market price and the historical price.*

Based on Figures 9 to 11, the using of the CRs for the uninformed traders are less (with a maximum close to 89.5) than the using by the informed traders (with a maximum close to 133.5). This result is consistent with Menkhoff et al. (2010) that the informed traders are highly sensitive to market conditions but are less so for the uninformed traders. If we compare A3 and BM, Figure 9 shows that the learning of the informed traders does not significantly affect the evolution process of the uninformed traders. We observe some small differences in CR1, CR4, CR17 and CR18. They indicate that, due to the learning of the informed traders, the uninformed traders care more about the relation between the bid-ask midpoint and the fundamental value (CR1), but less so about the relation between the market price and the historical price (CR4). Uninformed traders also care more about the imbalance of depth beyond the quote (CR17) and the imbalance between buy and sell (CR18). Intuitively, when the informed traders learn, they release more private information to the market so that the uninformed traders can benefit from the market imbalance between depth from the buy and the sell sides. For example, if informed traders find that the quote price is much higher than the fundamental value, they may submit more aggressive orders, which lead the depth of selling to become thinner; if informed traders find the quote price is slightly higher than the fundamental value, they may submit less aggressive limit buy orders, which leads the the depth of buying to become thicker. Therefore the market depth contains rich information from which the uninformed traders can benefit. This result also explain why the learning from uninformed traders significantly reduce the imbalance of the depth beyond the quote in Observation 9.

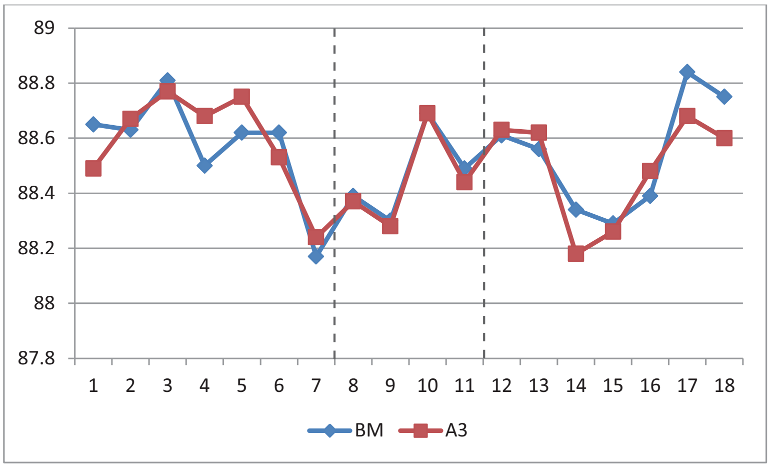

FIGURE 9. The average frequency of the classified rules CR1 to CR18 used by the uninformed traders per trader in experiments A3 and BM.

Observation 17. *When the market becomes more informative, the uninformed traders care more about the relation between the market price and the historical price, the sign of last transaction, and all the information related to the limit order book.*

This observation is based on the results presented in Figure 10 which compares B1 and B2. It shows that, when the proportion of the informed traders becomes extremely high, the uninformed traders care more about the market price and historical price (CR4) and the limit order book states (CR7 to CR18). Because of high informativeness, the historical price and the limit order states contain information released by the informed traders. While there is only one informed trader, the market becomes less informative. Then the uninformed traders care more about the information of the lag fundamental value (CR1, CR2, CR3 and CR5).

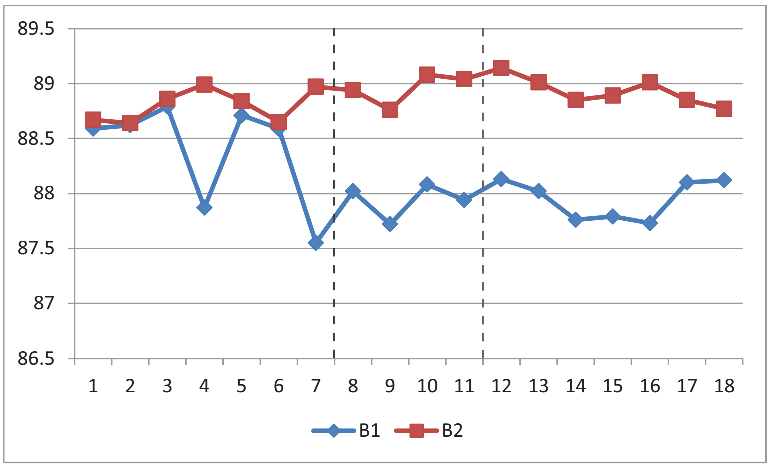

FIGURE 10. The average frequency of the classified rules CR1 to CR18 used by the uninformed traders per trader in experiments B1 and B2.

Observation 18. *When less weight is associated with the latest performance, the uninformed traders care less about all the market conditions.*

This observation is based on the results in Figure 11. It shows that all the experiments share a similar patter except the low use for E1. In this case, traders put less weight on the most recent performance. Therefore the uninformed traders do not pay much attention to the market conditions, which is intuitive.

Observation 19. *When information-lived time is shorter, the uninformed traders care more about all the market conditions.*

This observation is based on the results in Figure 11. It shows that uninformed traders care more about all the market conditions in D1 than D2. Intuitively, when the information-lived time is shorter, so the uninformed traders know the true fundamental value earlier and can update their performance more quickly, they should care more about the changing market conditions. This is also consistent with Observation 16, that is to say, in our model, the main difference between informed and uninformed is the asymmetric information, so when the information-lived time is shorter, the uninformed traders's evolution should be more close to the informed traders'.

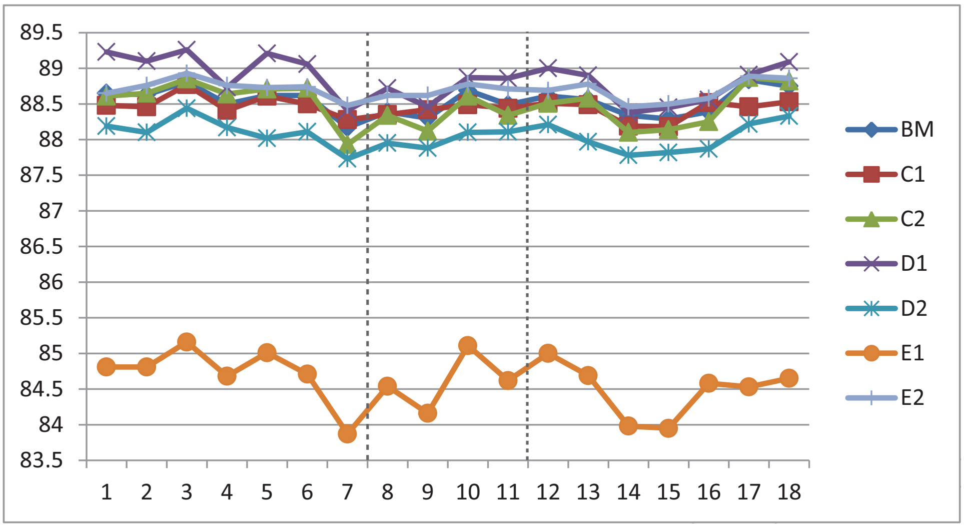

FIGURE 11. The average frequency of the classified rules CR1 to CR18 used by the uninformed traders per trader in experiments BM and Group C, D and E.

6. CONCLUSION

This paper proposes a two-sided learning model of an artificial limit order market with asymmetric information in which the trading strategies of the informed and uninformed traders evolve endogenously by genetic algorithm learning. We show that the model is able to capture many realistic features of limit order markets and provides some insight into the dynamics of limit order books.

When dealing with complicated financial markets, learning becomes very important and challenging. Learning of the uninformed traders affects not only themselves but also the informed traders and vice-verse. The limit order market model established in this paper demonstrates that genetic algorithm learning can be very effective in overcoming the challenge and facilitates the interaction of traders and market efficiency. We show that learning, in particular learning from the uninformed traders, can improve the market informational efficiency and lead the market price to converge to the fundamental value; learning can generate some realistic stylized facts in limit order markets; and learning can generates very rich limit order book phenomena. We also show that learning can affect order submission and market liquidity. Finally, we explore how the informed and uninformed traders process private and market information and how they interact with each other.

Order submission plays a key role in the understanding of the dynamics of the limit order book. We show that it is the learning from the uninformed instead of the informed traders that affects the order submission of both the informed and uninformed traders. Because of the learning from the uninformed traders, the informed traders tend to submit less aggressive limit orders but more market orders than the uninformed traders, while the uninformed traders use market orders to replace some aggressive limit orders. Overall the informed traders are more willing to consume liquidity than the uninformed traders, while their liquidity supply is about the same. However, the learning from the informed traders has less impact on order submission.

The model is able to generate a "hump" order book shape near the quote. We show that although the hump order book shape is generated by the continues double auction mechanism, nevertheless the learning of traders, in particular the uninformed traders, can reduce the order imbalance beyond the bid and ask. Learning from the informed traders widens the bid-ask spread and increases the spread volatility. In contrast learning from the uninformed traders narrows the bid-ask spread and reduces bid-ask spread volatility. Also the spread is mainly determined by the learning ability when price deviation from the fundamental value is large, but by the order aggressiveness when the price deviation is small. Due to the information advantage of the informed traders, their evolution of trading strategies can be very different from the uninformed traders. In general, the informed traders care more about the change in market price and depth at the quote, while the uninformed traders care less about the market conditions than the informed traders. When the market becomes more informative, informed traders care more about the depth away from the quote, while uninformed traders care more about the market price relative to the historical price and all the limit order book states. When traders put less weight on the most recent performance, uninformed traders tend to pay less attention to the market conditions. The learning can affect the evolution process differently. When uninformed traders learn, informed traders care about all the market conditions, especially the market price relative to the historical price and the sign of the last transaction. However learning from the informed traders has less impact on the evolution process of the uninformed traders, although they may care more about the bid-ask midpoint relative to the fundamental value.

Genetic algorithm learning has been used in economic literature. This paper demonstrates that it is also a very effective tool to study limit order markets. In particular, it can be used to model the learning of the uninformed traders and to explore its role on market efficiency and order book phenomena.

Appendix: The further analysis of order book shape

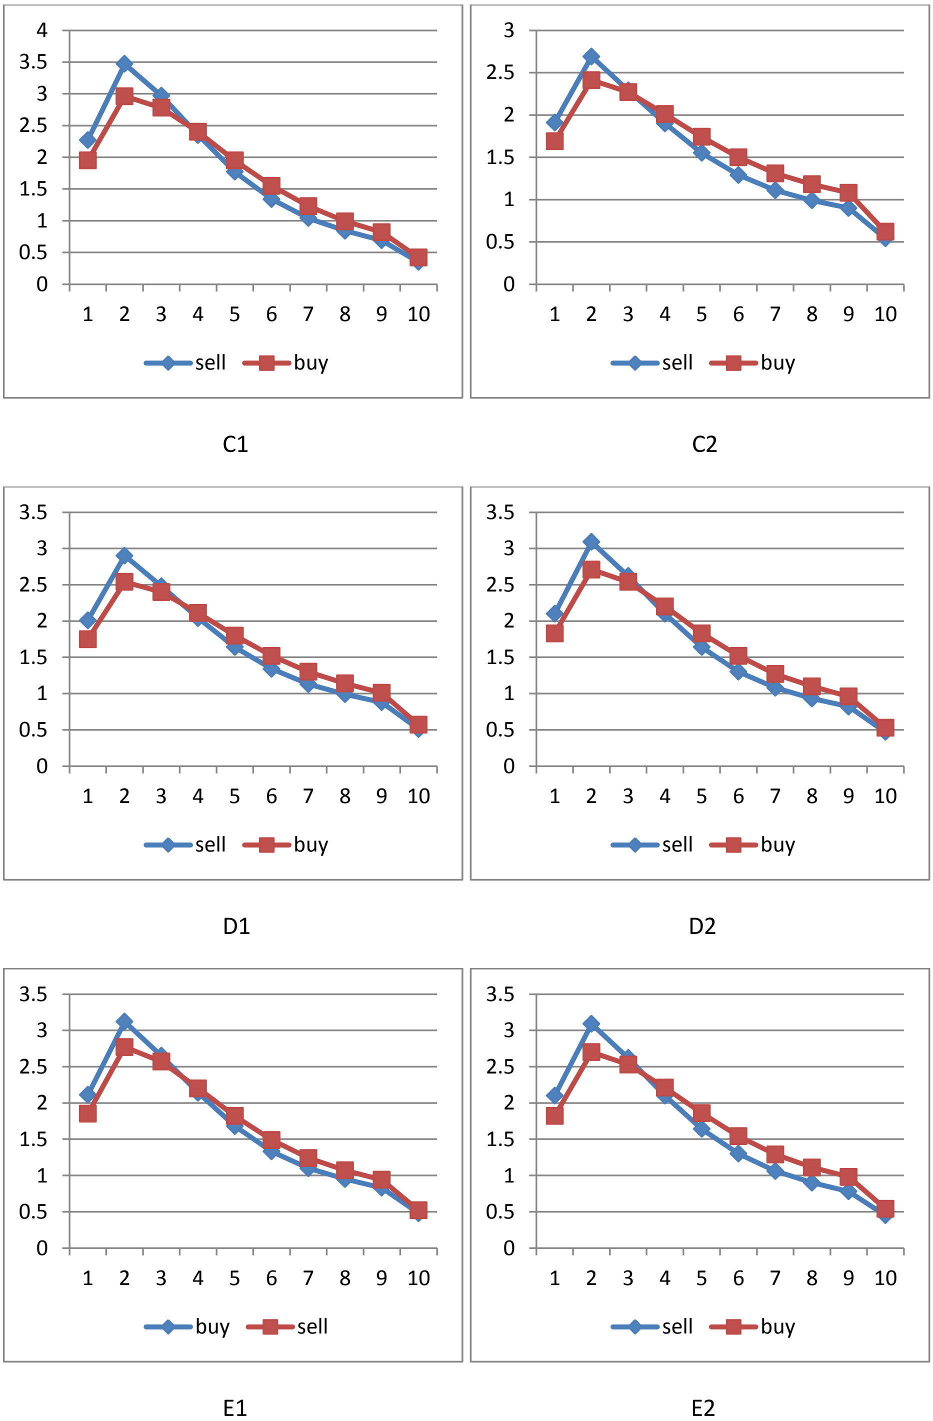

Figure 12. The order book shape in experiment Groups C, D and E.

References

Arifovic, J. (1994), 'Genetic algorithm learning and the cobweb model', *Journal of Economic Dynamics and Control* **18**, 3–28.

Arifovic, J. (1996), 'The behavior of the exchange rate in the genetic algorithm and experimental economies', *Journal of Political Economy* **104**, 510–541.

Arthur, W., Holland, H., LeBaron, B., Palmer, R. and Tayler, P. (1997), *The Economy As An Evolving Complex System II*, Addison-Wesley, chapter Asset pricing under endogenous expectations in an artificial stock market, pp. 15–44.

Bloomfield, R., O'Hara, M. and Saar, G. (2005), 'The 'make or take' decision in an electronic market: evidence on the evolution of liquidity', *Journal of Financial Economics* **75**(1), 165 – 199.

Chakravarty, S. and Holden, C. (1995), 'An integrated model of market and limit orders', *Journal of Financial Intermediation* **4**, 213?41.

Chen, S.-H. (2002), *Genetic algorithms and genetic programming in computational finance*, Kluwer, Boston,MA.

Chen, S.-H., Chang, C.-L. and Du, Y.-R. (2012), 'Agent-based economic models and econometrics', *Knowledge Engineering Review* **27**(2), 187–219.

Chiarella, C., Iori, G. and Perellò, J. (2009), 'The impact of heterogeneous trading rules on the limit order book and order flows', *Journal of Economic Dynamics and Control* **33**(3), 525 – 537.

Collin-Dufresne, P. and Fos, V. (2013), Do prices reveal the presence of informed trading?, Working paper, SSRN.

Dugast, J. (2012), Limit attention and news arrival in limit order markets, Working paper, HEC Paris.

Easley, D., Hvidkjaer, S. and O'Hara, M. (2010), 'Factoring information into returns', *Journal of Financial and Quantitative Analysis* **45**, 293?09.

Glosten, L. R. (1994), 'Is the electronic open limit order book inevitable?.', *Journal of Finance* **49**(4), 1127 – 1161.

Goettler, R. L., Parlour, C. A. and Rajan, U. (2005), 'Equilibrium in a dynamic limit order market', *Journal of Finance* **10**, 2149–2192.

Goettler, R. L., Parlour, C. A. and Rajan, U. (2009), 'Informed traders and limit order markets', *Journal of Financial Economics* **93(1)**, 67–87.

Gould, M., Porter, M., Williams, S., Fenn, D. and Howison, S. D. (2012), Limit order books, Working paper, Oxford University.

Harris, L. (1998), 'Optimal dynamic order submission strategies in some stylized trading problems', *Financial Markets, Institutions, and Instruments* **7**, 26?4.

Holland, J. H. (1975), *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence*, University of Michigan Press, Ann Arbor, MI.

Holland, J. H. and Miller, J. H. (1991), 'Artificial adaptive agents in economic theory', *American Economic Review* **81**(2), 365–370.

Hollifield, B., Miller, R. A., Sandas, P. and Slive, J. (2006), 'Estimating the gains from trade in limit-order markets', *The Journal of Finance* **61**(6), 2753–2804.

Kaniel, R. and Liu, H. (2006), 'What orders do informed traders used?', *Journal of Business* **79**, 1867–1913.

Keim, D. and Madhavan, A. (1995), 'Anatomy of the trading process: Empirical evidence on the behavior of institutional traders', *Journal of Financial Economics* **37**, 371?98.

Kluger, B. D. and McBride, M. E. (2011), 'Intraday trading patterns in an intelligent autonomous agent-based stock market', *Journal of Economic Behavior & Organization* **79**(3), 226 – 245.

Lin, H. and Ke, W. (2011), 'A computing bias in estimating the probability of informed trading', *Journal of Financial Markets* **14**, 625?40.

Menkhoff, L., Osler, C. L. and Schmeling, M. (2010), 'Limit-order submission strategies under asymmetric information', *Journal of Banking & Finance* **34**, 2665–2677.

Noe, T. H., Rebello, M. J. and Wang, J. (2003), 'Corporate financing: an artificial agent-based analysis', *Journal of Finance* **58**, 943–973.

Noe, T. H., Rebello, M. J. and Wang, J. (2006), 'The evolution of security designs', *Journal of Finance* **61**, 2103–2135.

O'Hara, M. (2001), Overview: Market structure issues in market liquidity, *in* 'Market liquidity: proceedings of a workshop held at the BIS', BIS Papers.

Parlour, C. A. and Seppi, D. J. (2008), *Limit order markets: a survey*, number Chapter 3 *in* 'Handbook of Financial Intermediation and Banking', North-Holland, pp. 63–96. in *Handbook of Financial Intermediation and Banking*, Eds. Boot, A.W.A. and Thakor,A.V.

Rosu, I. (2009), 'A dynamic model of the limit order book', *Review of Financial Studies* **22**, 4601–4641.

Rosu, I. (2012*a*), Liquidity and information in order driven markets, Working paper, SSRN.

Rosu, I. (2012*b*), *Order Choice and Information in Limit Order Markets*, number Chapter 2 *in* 'Market Microstructure: Confronting Many Viewpoints', Wiley, pp. 41–60. in *Market Microstructure: Confronting Many Viewpoints*, Eds. Abergel, F. and Bouchaud,J. and Foucault, T. and Lehalle, C.A. and Rosenbaum, M.

Routledge, B. (1999), 'Adaptive learning in financial markets', *Review of Financial Studies* **12**, 1165–1202.

Routledge, B. (2001), 'Genetic algorithm learning to choose and use information', *Macroeconomic Dynamics* **5**, 303–325.

Seppi, D. J. (1994), 'Liquidity provision with limit orders and a strategic specialist', *Review of Financial Studies* **10**, 103–150.

Theissen, E. (2000), 'Market structure, informational efficiency and liquidity: An experimental comparison of auction and dealer markets', *Journal of Financial Markets* **3(4)**, 333–363.

Wei, L., Zhang, W., He, X. and Zhang, Y. (2013*a*), An agent-based model of stylized facts and limit order book phenomena, Working paper, Quantitative Finance Research Central (QFRC), University of Technolgy, Sydney.

Wei, L., Zhang, W., He, X. and Zhang, Y. (2013*b*), Learning and information dissemination in limit order markets, Working paper (no. 333), Quantitative Finance Research Central (QFRC), University of Technolgy, Sydney.

Yan, Y. and Zhang, S. (2012), 'An improved estimation method and empirical properties of the probability of informed trading', *Journal of Banking & Finance* **36**(2), 454 – 467.