



WINPEC Working Paper Series No.E1514
October 2015

Dynamic predictor selection and order splitting in a limit order market

Ryuichi Yamamoto

Waseda INstitute of Political Economy
Waseda University
Tokyo, Japan

Dynamic predictor selection and order splitting in a limit order market

Ryuichi Yamamoto*
School of Political Science and Economics,
Waseda University,
Tokyo, 169-8050, Japan

This version: October 2015

Abstract

Recent empirical research has documented the clustered volatility and fat tails of return distribution in stock markets, and yet the markets are informationally efficient. Certain agent-based theoretical models attempt to explain the empirical features in terms of investors' order-splitting or dynamic switching strategies, both of which are popularly used by actual stock investors. However, little theoretical research has discriminated the behavioral assumptions within a model and compared the impacts of the assumptions on the empirical features. In addition, the research has not simultaneously replicated the return features and empirical features on market microstructure, such as patterns of order choice. This study constructs an artificial limit order market, in which investors split orders into small pieces or interchangeably use fundamental and trend-following predictors over time. We demonstrate that, on the one hand, the market that features strategies with order splitting and dynamic predictor selection can independently replicate clustered volatility and fat tails in an informationally efficient market. However, we also show that patterns of order choice do not match those found in certain previous empirical studies in both types of economies. Thus, we conclude that in reality, the two strategies simultaneously work to generate the empirical macroeconomic features, but that investors may also use certain other strategies in actual stock markets. In addition, we demonstrate that the impact of both strategies on the volatility persistence tends to be greater as the number of traders increases in the market; this finding implies that the order-splitting strategy and dynamic predictor selection are more crucial on the empirical phenomena that pertain to larger capital stocks.

JEL classifications: G12; G14

Keywords: Dynamic predictor selection; Order splitting; Volatility clustering; Fat tails; Information efficiency; Limit order market; Agent-based modeling

* Corresponding author: Tel.: +81 3 3208 0534; fax: +81 3 3204 8957.
E-mail address: ryuichi@waseda.jp (Ryuichi Yamamoto).

1. Introduction

Agent-based modeling is a useful computational tool in crafting economic models to understand the relation between dynamically updating trading strategies and stylized empirical features in stock markets. Certain empirical facts in stock markets, such as the clustered volatility and fat tails of return distribution, have been successfully explained by previous agent-based theoretical models that make use of investors' order-splitting and dynamic switching strategies, both of which are popularly used by actual stock investors.¹ However, little theoretical research has discriminated the behavioral assumptions *within a model* and compared the impacts of the assumptions on the empirical features. Moreover, few theoretical works have simultaneously replicated certain empirical macroeconomic facts, such as clustered volatility and market microstructure features, such as order placement patterns. This study constructs an artificial limit order market, in which stock investors adopt order-splitting and dynamic switching strategies. We hypothesize that traders' order-splitting and switching strategies cause clustered volatility and fat tails in an informationally efficient market.² This paper thus undertakes computer-based experiments on our market that directly support the presented hypothesis. However, we show that neither behavioral assumption is related to the generation of the empirical fact on the order placement patterns. We also demonstrate that order splitting and dynamic switching strategies tend to intensify the volatility persistence as the number of traders in the market grows. The

¹ For example, Hommes (2006) and LeBaron (2006) summarize the theoretical success of agent-based models. Order splitting is a trading strategy commonly used in real stock exchanges. According to the estimation by Vaglica, Lillo, Moro, and Mantegna (2008), large orders are divided into smaller pieces of possibly 1,000 to 10,000 units. The survey studies by Lui and Mole (1998), Menkhoff and Taylor (2007), and Menkhoff (2010) provide evidence of the widespread use of fundamental and trend-following strategies by financial market participants. Brock and Hommes (1997, 1998) and Lux (1995), meanwhile, theoretically model dynamic strategy selections, and the strategy is empirically validated by several other studies, such as that of Boswijk, Hommes, and Manzan (2007).

² For example, Mandelbrot(1972) and Ding, Granger, and Engle (1993) demonstrate volatility clustering, while Farmer and Lillo (2004) and Lux (2001) find the fat tails. Bouchaud, Gefen, Potters, and Wyart (2004) and Lillo and Farmer (2004) each show the information efficiency in stock markets.

result implies that within real stock markets, the importance of both strategy types on the empirical phenomena increases among larger capital stocks.

We model a continuous double-auction or limit order market, one of the modern and standard market institutions used in actual stock exchanges. Our agents first choose the total order size by maximizing the expected utility. We then allow our agents to follow an order-splitting or switching strategy. When investors follow an order-splitting strategy, they observe the order-book condition and split the order by minimizing the expected market impact of a trade. Our agents dynamically adjust the size of their splits, while the order-book condition changes in the following manner. On the one hand, when the order book is being filled up with limit orders, our agents place their splits possibly at a large size, because the price impact of a trade would thus be smaller. On the other hand, they execute smaller orders when the liquidity in the order book declines. When our investors take the switching strategy, they dynamically select a fundamental or trend-following strategy that has generated higher returns in the recent past, as in the model originally developed by Brock and Hommes (1997, 1998). The setup is consistent with the concept of myopic utility maximization.

We simulate our market with an order-splitting strategy independent of the market featuring the dynamic switching rule. We evaluate the significance of the behavioral assumptions for the generation of not only clustered volatility and fat tails in an informationally efficient market, but also the empirical properties of the order submissions namely, 1) the distribution of the number of split orders and 2) the absolute and conditional frequencies of order choice.³ The behavioral assumptions which pertain to order-splitting and/or dynamic switching strategies, can

³ We introduce in Subsection 3.1 details on these phenomena.

be significant in actual stock markets, if all the empirical properties on returns and order placements are successfully replicated.

The absolute and conditional frequencies of order submissions have been empirically documented in several previous empirical studies, including some on the Paris Bourse, the Toronto Stock Exchange, and the Australian Stock Exchange. On the one hand, the order ratio (absolute frequency) of market order submissions to all orders placed to the order book is approximately 28–53%, while that of limit order placements outside the spread is 27–65%. That of limit orders inside the spread, meanwhile, is approximately 2–19%. On the other hand, the results of past studies have also shown the conditional frequency of order types with regard to depth and spread. Stock investors submit more aggressive orders, such as market orders, as the depth on the same side of the order book becomes thick or the spread narrows, and they place less aggressive orders, such as limit orders outside the spread, as the depth thins or the spread widens.⁴ We evaluate the significance of our behavioral assumptions in the market by investigating whether the order patterns are also consistent with the empirical evidence.

We demonstrate that economies that follow an order-spitting or dynamic switching strategy can generate clustered volatility and fat tails in an informationally efficient market. However, an order-splitting economy produces a much higher frequency of market order submissions and a much lower frequency of limit order provisions than those seen in empirical evidence from actual stock markets. In addition, the market persistently receives market orders even when the depth thins or the spread widens. In a market with dynamic predictor selection, we find that order choice is not influenced, regardless of depth condition. Therefore, the results

⁴ In the literature, the most aggressive order is the market order; meanwhile, limit orders within the spread are more aggressive than limit orders outside the spread, but are less aggressive than market orders. See, for example, Biais, Hillion, and Spatt (1995), Duong, Kalev, and Krishnamurti (2009), Griffiths, Smith, Turnbull, and White (2000), Hall and Hautsch (2006), Handa, Schwartz, and Tiwari (2003), and Rinaldo (2004).

suggest that order-splitting and trend-following strategies can, quite independently, be the cause of clustered volatility and fat tails in an informationally efficient market. However, we may need to make certain other behavioral assumptions to match empirical properties to the order submissions observed in actual stock markets. The behavioral assumption could be a strategy that chooses limit and market orders interchangeably, depending on the demand–supply imbalance in the order book.

This study makes the following five contributions. First and most significantly, our agent-based model discriminates the behavioral assumptions within a model and evaluates the impacts of the assumptions on the return distribution and order patterns. Previous studies explain clustered volatility and/or fat tails by imposing a behavioral assumption in the model, such as herding (e.g., Alfarano, Lux, and Wagner, 2008), order splitting (e.g., Lillo, Mike, and Farmer, 2005), the trend-following strategy (e.g., Chiarella, Iori, and Perelló, 2009), or dynamic predictor selection (e.g., Lux, 1995). However, none of the previous studies models order splitting and dynamic predictor selection simultaneously within one model. Thus, we are not able to argue which behavioral assumption is the more important cause of clustered volatility, fat tails, and the empirical features of order patterns; such is the subject of this study. Furthermore, as pointed out by Chen, Chang, and Du (2012), Chiarella, He, and Wei (2015), and Gould, Porter, Williams, Fenn, and Howison (2013), certain progress has been made in agent-based modeling for replicating empirical macroeconomic phenomena such as clustered volatility. However, there has been limited progress in pinpointing the causes of stylized facts in the market microstructure, such as order placement patterns, and thus this represents the next mission of agent-based modeling in the financial market. The current study contributes to the body of the knowledge on this issue by relating the two behavioral assumptions to the empirical order choice phenomena.

The second contribution made by the current study is that we model a limit order market by relaxing the restrictive assumptions otherwise imposed in previous theoretical studies on limit order markets; additionally, we investigate the relation between the more realistic setup and the empirical features of returns and order patterns. The studies of Goettler, Parlour, and Rajan (2005), Foucault, Kadan, and Kandel (2005), and Roşu (2009) model agents' strategy choices between limit and market orders, as well as the equilibrium order placement strategies in the limit order market. Finding the equilibrium and the conditions underpinning that equilibrium in the limit order market is important for understanding the fundamental features of the market structure. Meanwhile, they impose restrictive assumptions on their analyses. For example, in Foucault, Kadan, and Kandel (2005), agents cannot observe the order-book conditions and thus the strategy is independent of the state of the limit order book; Roşu (2009), meanwhile, assumes that agents can trade only one unit at a time. In Goettler, Parlour, and Rajan (2005), investors do not consider the dynamic switching strategy and the cost-minimizing means of splitting the order. This study relaxes these assumptions, and furthermore, considers the popularly utilized investment strategies in actual limit order markets namely, optimal order-splitting and switching strategies, in order to investigate the empirical features of a high-frequency time series.

Our third contribution is in terms of our modeling of the order-splitting strategy. Our agents endogenously and dynamically determine the size of sliced orders by observing the order-book condition that is adjusted over time in response to the actual transaction history. Order splitting has been modeled by several approaches in the literature. Our model is complementary to but unique from previous works in the following three ways. First, previous studies have usually assumed that the total order size is exogenously given and that a series of split order sizes executed in the future is determined at the initialization of their trades. These models have

focused exclusively on the endogenous determination of split sizes by minimizing the implementation shortfall namely, the difference between the price (or value) when agents initialize the trade and the execution price (or value).⁵ The agents in our model, in contrast, endogenously determine the total order size by maximizing their utility. In addition, they decide the optimal sizes of splits by minimizing the implementation shortfall; however, the split sizes are dynamically adjusted in response to the prevailing order-book conditions. Second, as in our model, the agents in the models by Alfonsi, Fruth, and Schied (2010) and Obizhaeva and Wang (2005) optimally select the split size based on the order-book conditions. The previous models assume the time-invariant distribution of the order book. However, our model allows the order-book shape to change over time, as in real stock exchanges, in response to the transaction history. Thus, our agents decide the split size according to the most recent order-book information. Third, as in the models presented by Almgren and Lorenz (2007), Bertsimas and Lo (1998), and Lorenz and Almgren (2011), our agents dynamically update their split order sizes according to the market impact of the trades; however, those previous works assume that the functional form of price impacts are time-invariant. In our economy, in contrast, the market impact depends on current order-book conditions that change over time in response to the execution, submission, and cancellation of limit orders.⁶

Fourth, this study provides an order-splitting model that reproduces the power-law distribution of the number of splits, which is consistent with the empirical evidence presented by Vaglica, Lillo, Moro, and Mantegna (2008). Our agents divide their trades into multiple transactions in order to mitigate market impact risk. They thus select the size of the splits by

⁵ See, for example, Alfonsi, Fruth, and Schied (2010), Almgren and Chriss (2000), Almgren and Lorenz (2007), Lorenz and Almgren (2011), and Obizhaeva and Wang (2005).

⁶ In addition to the previous models, order splitting is modeled by Lillo, Mike, and Farmer (2005) and Yamamoto and LeBaron (2010). However, the total order size and split size are randomly determined in those models.

observing the order-book conditions and computing the likely market impact. In this way, we theoretically explain why order-splitting investors make a large number of splits in actual markets.

Fifth, we formulate an agent-based model where agents' actions are determined by both price expectation and order-book information. Agent-based models usually assume that strategies are formed on the basis of price expectations.⁷ However, in real stock exchanges, order-book information is revealed to the public in real time and stock investors in reality utilize it to form their strategies, because that information can provide ideas on future market dynamics. Thus, traders' strategies in agent-based models should reflect information not only on past prices, but also that from the order book.

The remainder of this paper is organized as follows. Section 2 describes our market structure. Section 3 presents the results from the computer simulations on the economies, which make use of order splitting and dynamic predictor selection. With regard to each strategy, we first analyze clustered volatility, fat tails, and information efficiency; we next investigate the distribution of the number of splits, the absolute frequency of order types, and the conditional frequency of order types on the depth and spread in our market. A brief summary and suggestions for future research conclude the article in Section 4.

2. The model

This section designs our limit order market, which is a continuous double auction or order-driven economy. In this economy, agents can place market orders that are executed

⁷ Among the several agent-based models of this type, see, for example, Brock and Hommes (1998), Chiarella, Iori, and Perelló (2009), Frankel and Froot (1990), Kirman (1993), and LeBaron, Arthur, and Palmer (1999). Very few agent-based models incorporate order-book information into agent' trading decisions. See, for example, Anufriev, Arifovic, Ledyard, and Panchenko (2011) and Yamamoto (2011).

immediately at the best available prices. They can also place limit orders that specify the size and the limit price at which they are willing to trade, and their limit order (or a part of the order) is executed at their pre-specified or better price when other agents hit the order. There is no stock transaction intermediary in our clearing mechanism, such as a dealer, specialist, or designated market maker. The market mimics modern electronic markets with an open order book, such as that on the Tokyo Stock Exchange after the opening session and before the closing session. This simplified setup prevents us answering institutional questions in a market in which dealers or specialists play a role; however, it does highlight how order splitting and dynamic adopting strategies affect a market. In the following, we first determine the total size of the orders by modifying the setups in Chiarella, Iori, and Perelló (2009). Our agents form expectations and optimally choose the total order size by maximizing the myopic expected utility. Then, we illustrate our order-splitting and dynamic switching mechanisms, in order.

In our market, N agents that are indexed to i trade a risky asset. We define p_t as the price of the risky asset at t that will be determined by each trade, or the average of the best bid and ask prices when there are no transactions at time t . We assume that agents can observe all past prices.

The utility function is expressed in a form of constant absolute risk aversion (CARA), as:

$$\hat{E}_t^i(U(w_{t+L_i}^i)) = \max \hat{E}_t^i(-\exp(-\gamma w_{t+L_i}^i)), \quad (1)$$

where \hat{E}_t^i denotes the best forecast of agent i at time t and w_t^i is wealth at t . γ is a CARA coefficient. The agents' budget constraint is given as:

$$w_{t+L_i}^i = w_t^i + d_t^i p_{t-1} \ln(\hat{p}_{t,t+L_i}^i / p_{t-1}), \quad (2)$$

where $\hat{p}_{t,t+L_i}^i$ is the L_i period-ahead prediction of prices by agent i at t . By assuming a Gaussian distribution for the return series, we can rewrite the expected utility function as:

$$\hat{E}_t^i(-\exp(-\gamma w_{t+L_i})) = -\exp(-\gamma \hat{E}_t^i(w_{t+L_i}) + \gamma^2 V_t^i(w_{t+L_i})/2), \quad (3)$$

where $V_t^i(\bullet)$ refers to the conditional variance estimate of returns by agent i , and is assumed to be constant. The moments in Equation (3) can be re-expressed as:

$$\begin{aligned} \hat{E}_t^i(w_{t+L_i}) &= w_t + d_t^i p_{t-1} \ln(\hat{p}_{t,t+L_i}^i / p_{t-1}) \\ V_t^i(w_{t+L_i}) &= d_t^{i2} p_{t-1}^2 V_t^i(\ln(\hat{p}_{t,t+L_i}^i / p_{t-1})). \end{aligned}$$

Thus, the utility function becomes:

$$\begin{aligned} &\hat{E}_t^i(-\exp(-\gamma w_{t+L_i})) \\ &= -\exp(-\gamma w_t) \exp\left(-\gamma d_t^i p_{t-1} \ln(\hat{p}_{t,t+L_i}^i / p_{t-1}) + \gamma^2 d_t^{i2} p_{t-1}^2 V_t^i(\ln(\hat{p}_{t,t+L_i}^i / p_{t-1}))/2\right). \end{aligned} \quad (4)$$

Agent i maximizes Equation (4) in terms of d_t^i and sets the myopic optimal demand at time t according to:

$$d_t^i = \frac{\ln(\hat{p}_{t,t+L_i}^i / p_{t-1})}{\gamma V_t^i p_{t-1}}. \quad (5)$$

The L_i period-ahead prediction of returns by agent i at t is defined as $\hat{r}_{t,t+L_i}^i$, while in the following subsections we detail the expectation-formation mechanisms in order-splitting, no order-splitting, and dynamic switching strategies. Based on the return forecast and current price, the L_i period-ahead prediction of prices by agent i at t is given as follows:

$$\hat{p}_{t,t+L_i}^i = p_{t-1} e^{\hat{r}_{t,t+L_i}^i}. \quad (6)$$

The order size for agent i at time t , denoted by s_t^i , is determined by the absolute difference in asset demands at t and $t-1$:

$$s_t^i = \text{abs}(d_t^i - d_{t-1}^i). \quad (7)$$

When the difference, $d_t^i - d_{t-1}^i$, is positive (negative), an agent i buys (sells) shares by s_t^i . Following the approach proposed by Anufriev and Panchenko (2009) and Mike and Farmer's (2008) empirical finding about order placement distribution, agent i sets his or her limit price at t according to a random draw such that the limit price relative to the best price at t follows a Student's t -distribution with 1.3 degrees of freedom. The best bid and best ask at $t-1$ are utilized as the best prices for a buy order and a sell order, respectively. We denote the limit price of a buy order as b_t^i and that of a sell order as a_t^i , while the limit prices are independent over time and across agents.

When b_t^i is below the best ask or a_t^i is above the best bid, agent i places a limit order of the total amount s_t^i at the requested price a_t^i or b_t^i . The agent places a market order when the limit price crosses the best quotes. When the trade takes place, the matched limit order at the best price is removed from the book. When all limit buy orders at the prevailing best bid are removed from the book, the best bid moves down by one tick. Similarly, the best ask moves up by one tick when all the limit sell orders at the best ask are removed from the book. t is the smallest time step (or a click) that represents one action (placing a limit or market order) by an agent. We denote that tick size as Δ , and a limit order expires in τ actions after the placement.

In the following subsections, we first describe the order-splitting mechanism and then the switching mechanism in details.

2.1. Order-splitting strategy

In our economies with and without an order-splitting strategy, we assume that a noise-trading component determines agents' expectations. This assumption is made to highlight the

impact of the order-splitting strategy on the market while isolating the impacts from any other forms of expectation rules, such as a trend-following rule. The expectation rule is illustrated as:

$$\hat{r}_{t,t+L_i}^i = n^i \varepsilon_t, \quad (8)$$

where n^i is initially given for each agent i from $n^i \sim |N(0, \sigma_n)|$ and $\varepsilon_t^i \sim N(0,1)$.

In our order-splitting economy, all agents are assumed to slice their orders into pieces only when placing market orders.⁸ Our order-splitting agents divide the total order s_t^i into smaller units and execute them piece by piece in order to minimize the expected transaction cost from the market impact of a trade. Thus, in our setup, the split size is endogenously determined. We assume that all limit orders on both sides of the book and order volumes at those prices are observable prior to trading and are consistent with several actual stock exchanges, such as the Australian Stock Exchange, the London Stock Exchange, the NASDAQ, NYSE OpenBook, and the Tokyo Stock Exchange. Our agents observe this information at t and calculate a likely deviation of the expected executed value from the value evaluated at the price they initialize the trade as:

$$E_t[C] = |x_t^i S_0 - x_t^i \hat{S}_{t+1}|, \quad (9)$$

where S_0 is the security price immediately before the agent starts executing the orders. This is called the decision or arrival price, which is the price of the stock that prompted the decision to trade. In reality, the broker, who receives the orders in order to minimize the implementation

⁸ Actual investors also split their limit orders as hidden or iceberg orders (e.g., Bessembinder, Panayides, and Venkataraman, 2009), or place their slices across different trading platforms. However, we focus only on the slicing strategy of market orders for the following two reasons. First, we can compare our results on the empirical distribution of the number of market-order splits to those in reality. In this manner, we can also investigate how the distribution of the number of market-order splits generated by our model is related to the empirical return distributions and order patterns in reality. Second, we can also compare our model to those in the previous order-splitting models, as presented in the Introduction.

shortfall, normally uses the last transaction price or the mid-quote. We assume that our agents utilize the last recorded prices (i.e., the transaction prices when the last orders were executed and mid-prices if not executed) for S_0 . x_t^i is the split size, and \hat{S}_{t+1} is the expected execution price averaged by the shares traded at $t+1$. Our agents refer to the order book depths to forecast the expected execution price. For example, suppose that the current best ask is \$100 and the next best ask is \$110, and that there are 150 shares of limit orders at the best ask and 220 shares at the next best ask. If the agent plans to buy 200 shares immediately namely, $x_t^i = 200$, the expected average execution price \hat{S}_{t+1} is $(\frac{150}{200}) \times \$100 + (\frac{50}{200}) \times \$110 = \$102.5$. Therefore, $E_t[C]$ is computed as: $|200 \times \$100 - 200 \times \$102.5| = 500$.

This setup is comparable to those of Alfonsi, Fruth, and Schied (2010) and Obizhaeva and Wang (2005). As in our setup, the expected average execution price in these two studies is determined on the basis of the order-book conditions. Alfonsi, Fruth, and Schied (2010) assume that the functional form of the order book is constant over time, while Obizhaeva and Wang (2005) present a model with a constant order-book depth. However, our setup allows not only the order-book depth to vary over time as agents actually trade shares through the order book but also agents to make trading decisions by referring to prevailing order-book conditions.

Agent i endogenously determines the split size at t by minimizing the expected transaction cost. Agent i submits the market order according to the size that satisfies the condition:

$$E_t[C] = |x_t^i S_0 - x_t^i \hat{S}_{t+1}| = R^i, \quad (10)$$

where $R^i \in [0, R_{\max}]$. We assume that R^i is randomly assigned initially, but becomes constant over time. We set R_{\max} to equal 1. R^i measures trader i 's patience in accepting a certain

deviation of the expected execution price from the decision price. The setup means that agents allow themselves a small permissible error in expectations, but attempt to reduce the market impact risk to as small a level as possible.

Equation (10) indicates that agents evaluate the cost of execution at every transaction, and thus that the strategy is dynamically updated over time so as to minimize the expected transaction cost. The equation suggests that when the market is sufficiently liquid in that the order book contains a large number of limit orders, the split size x_t^i will be large as \hat{S}_{t+1} will not deviate markedly from the initial price S_0 . In contrast, the split size tends to be small, as the order book contains less liquidity. After executing his or her split, the agent computes the remaining order size z_t^i .

Previous studies on optimal order-splitting strategies, such as Alfonsi, Fruth, and Schied (2010), Almgren and Chriss (2000), Almgren and Lorenz (2007), Lorenz and Almgren (2011), and Obizhaeva and Wang (2005), assume an exogenously given time to complete all splits. However, in our model, the total time to execute all splits is endogenously determined, because the total time depends on the order-book condition, the total size of the orders, and the split size. The total time to complete the execution of the order will be shorter, when the order book becomes thick, agents have a small total order size, and they execute larger splits each time.

Our setup also indicates that the size of the market order becomes larger as the depth on the opposite side becomes thicker; this implies that the other side of the order book is important in determining an agent's order choice. This is consistent with the results of the theoretical work conducted by Parlour (1998) and the empirical findings of Duong, Kalev, and Krishnamurti (2009), Griffiths, Smith, Turnbull, and White (2000), and Ranaldo (2004).

Our order-splitting strategy is closely related to the arrival price strategy or implementation shortfall strategy that is heavily used in reality, not only in the electronic trading of equities but also in executing large orders. In the arrival price strategy or implementation shortfall algorithm, investors select the pre-trade price as the execution benchmark and execute the split orders so as to minimize the difference between the pre-trade benchmark and the execution price.⁹ The difference derives entirely from the market impact of a trade. The strategy indicates that investors face a trade-off between executing slowly versus rapidly. On the one hand, this difference would be smaller if investors were to trade as slowly as possible by executing small units each time. On the other hand, it would be larger if the price volatility were very large. In this case, investors should execute their orders quickly. The trade-off in the arrival price strategy or implementation shortfall algorithm suggests that the split size and thus the time to execute all splits depends on the shape of the order book, as outlined in the model of this study.

The timing for agent i to enter the market is determined as follows. Once any agents start executing splits, we assume that they try to place their splits in succession so as to mitigate future market impact risk. This implies that they aim to mitigate the future market impact risk that arises from delaying the execution of their splits. The assumption is consistent with the empirical finding of Tóth, Palit, Lillo, and Farmer (2011), who show the persistence of order signs placed by single investors (i.e., investors in the same exchange membership identifier in their sample). They show that autocorrelations tend to be larger as event time lags shorten, indicating that investors tend to execute their splits sooner. We accordingly set the probability that agent i , who

⁹ For example, according to Sussman (2005), the arrival price and implementation shortfall algorithms are used in 32% of algorithmic trading. The other popular strategy that minimizes the difference between execution prices and the benchmark price is the so-called Volume Weighted Average Price (VWAP) algorithm where investors trade such that the average execution price approaches a benchmark price (i.e., the VWAP), over a specific time interval. In algorithmic trading, the share of VWAP trade is 27%. See, for example, Perold (1988) and Lorenz and Almgren (2011) for further discussion of the implementation shortfall.

has already entered the market but has remaining orders to be sliced, places his or her split at t as:

$$prob_t^i = \frac{z_t^i}{NSP + \sum_{i=1}^N z_t^i}, \quad (11)$$

where z_t^i is the size of the orders remaining at t for agent i , who has already started splitting the orders. $\sum_{i=1}^N z_t^i$ is the total size of the orders waiting to be split at t . NSP is a parameter that relates to the probability that a new agent i will place his or her order. Thus, the probability of a new agent i entering the market is given by:

$$prob_t^i = \frac{NSP}{NSP + \sum_{i=1}^N z_t^i}, \quad (12)$$

where we set NSP equal to 1. Thus, the probabilities that a new agent will enter the market and that order-splitting agents will place their splits sum to 1. This setup indicates that it is less likely for a new agent to place and/or execute his or her order at time t if a larger number of agents or a larger size of orders is waiting to be split at t . The mechanism generates persistence in the order flow coming from a given agent, in that the order flow is less likely to shift to a new agent when more orders are waiting to be executed. This means that order executions are more condensed within certain time periods as they start splitting their orders, which is consistent with the empirical finding by Tóth, Palit, Lillo, and Farmer (2011). When a new agent takes an action, we randomly select the agent from the pool of agents that have not initiated order splitting.

2.2. Switching strategy

This subsection presents the model of dynamic predictor selection by following the setup by Brock and Hommes (1998). In the context of this setup, all agents' expectations are formed

by fundamentalist or trend-following indicators; this is consistent with the results of certain survey studies of financial market participants, such as those in Lui and Mole (1998), Menkhoff and Taylor (2007), and Menkhoff (2010).

We define the return at time t as r_t , and the mean return over the past L as:

$$r_t = \ln\left(\frac{p_t}{p_{t-1}}\right). \quad (13)$$

and

$$\bar{r}_L = \frac{1}{L} \sum_{j=1}^{L-1} \ln\left(\frac{p_{t-j}}{p_{t-j-1}}\right). \quad (14)$$

The fundamental and trend-following predictors are given respectively as:

$$\hat{r}_{t,t+L}^i = g_1 \ln\left(\frac{p^f}{p_{t-1}}\right) \quad (15)$$

$$\hat{r}_{t,t+L}^i = g_2 \bar{r}_L, \quad (16)$$

where p^f is the fundamental price of the risky asset. We assume that all agents can calculate the price, but need to pay the cost by C_F . g_1 and g_2 are the respective weights of the fundamentalist and trend-following predictors for agent i ; g_2 can take only positive values. In the dynamic predictor selection model by Brock and Hommes (1998), agents are assumed to select one of the two rules, meaning that there are only two rules in the economy. Thus, we assume that the values for g_1 , g_2 , and L are the same across agents.

In our economy with the switching strategy, an agent is randomly selected at t and conducts an action (i.e., placing a market or limit order). After observing S actions by agents, we assume that agents look back on the trading performances from the two rules for the past S

periods and update their predictors every S periods. We call S the learning frequency. Our agents select the trend-following predictor at t according to the following logit probability as:

$$n_{TF,t} = \frac{\exp\left(\beta \times \sum_{j=0}^{S-1} (-fitness_{TF,t-S+j})\right)}{\exp\left(\beta \times \sum_{j=1}^{S-1} (-fitness_{TF,t-S+j})\right) + \exp\left(\beta \times \left(\sum_{j=1}^S (-fitness_{F,t-S+j})\right) - C_F\right)}, \quad (17)$$

where β is the intensity of choice parameter that describes the tendency to switch the predictor.¹⁰ A larger value of β allows agents to switch the rule more frequently, while a lower value of β increases the likelihood of retaining the use of the same rule. We define $fitness_{F,t-S+j}$ and $fitness_{TF,t-S+j}$ as the squared forecast errors at $t-S+j$ according to the fundamental and trend-following rules, respectively. $(1 - n_{TF,t})$ is the likelihood that agents will utilize the fundamental predictor at t .

3. Model experiments

This section first introduces evidence of clustered volatility and fat tails in an informationally efficient market as well as the order patterns found in the previous empirical studies. We then simulate an economy without the order-splitting and switching mechanisms (a benchmark economy) in order to highlight the impacts of order-splitting and dynamic switching strategies on the return features and order patterns. We demonstrate that none of the empirical features are generated in the benchmark economy. We then independently examine order-splitting and switching economies, demonstrating that our model featuring the order-splitting and switching strategies can explain volatility clustering and fat tails in an informationally efficient

¹⁰ As introduced in detail by Brock and Hommes (1998), this probability is derived by solving the maximization problem by $\max\left(\sum_{j=0}^{S-1} (-fitness_{F,t-S+j}) + \omega_t, \sum_{j=0}^{S-1} (-fitness_{TF,t-S+j}) + \varpi_t\right)$, where ω_t and ϖ_t are random variables that are independent and extreme-value distributed

market. However, we will show that the behavioral assumptions may not be direct causes of all the order placement patterns. In the process, we also investigate how our results vary with the market size by changing the number of traders in the market.

In our simulations, one click represents an agent's action, such as a market order submission or a limit order provision. For our analyses on return features, we convert our clicks into wall-clock time and examine the time series. Specifically, we consider 25 actions as representing one wall-clock time. Using recorded prices (i.e., transaction prices when orders are executed and mid-prices if they are not), returns are computed as the percentage change in prices over 25 actions. We measure volatility by taking the standard deviations of the log returns over 25 actions. We simulate our economy for 60,000 wall-clock times in each run and use the final 10,000 clock times for the analysis of clustered volatility and fat tails. We simulate our economy using the parameter values given in Table 1. N , the number of agents, is set to 1,000 for our benchmark simulation. Later, we change this value to examine how the market size influences our results.

Table 1: Parameters

Number of agents: N	1,000
CARA coefficient and conditional variance: $\gamma\mathcal{V}$	0.01
Time horizon in chartist component: L	500
Tick size: Δ	0.01
Order life: τ	20,000
Std. of noise trading component: σ_n	5
Parameter of the fundamentalist component: g_1	10
Fundamental value: p^f	1,000
Cost for fundamental information: C_F	0.1
Parameter of the trend-following component: g_2	10
Learning horizon: S	250
Intensity of choice: β	10

3.1. Stylized facts and simulation results in a benchmark economy

3.1.1. Empirical return features

Table 2 presents the summary statistics of the return distribution of the S&P500 index, as well as the results from our benchmark economy where agents follow neither order-splitting nor dynamic switching rules. In the following tables, the bold numbers are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs. The first two rows in Table 2 summarize the descriptive statistics of daily returns for the S&P500 from January 2, 1990 to April 18, 2012, and for economies without order-splitting and dynamic switching strategies.¹¹ The returns in our economy without order-splitting and predictor selection strategies show smaller variations than those of the S&P500. For example, with S&P500 index, the standard deviation is 1.134/100 and kurtosis is 11.81; in the economy without order-splitting and dynamic predictor selection strategies, they are 0.048/100 and 3.582, respectively. The kurtosis in our benchmark economy shows that the return distribution has slightly fatter tails than those indicated under normal distribution. This possibly arises from the fact that we compute the returns from transaction prices rather than solely mid-quotes. Transactions are conducted at the best ask and bid prices, so as to generate greater variations of returns than those through the mid-quote returns.

Table 2: Summary statistics of returns of S&P500 and economies without order splitting and dynamic predictor selection. The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs.

	Mean $\times 10,000$	Std. $\times 100$	Skewness	Kurtosis
S&P500	2.719	1.134	-0.240	11.81
No order splitting and switching	0.005 (0.022)	0.048 (0.001)	-0.002 (0.028)	3.582 (0.054)

¹¹ We use this index from Datastream. The sample size is 6,451.

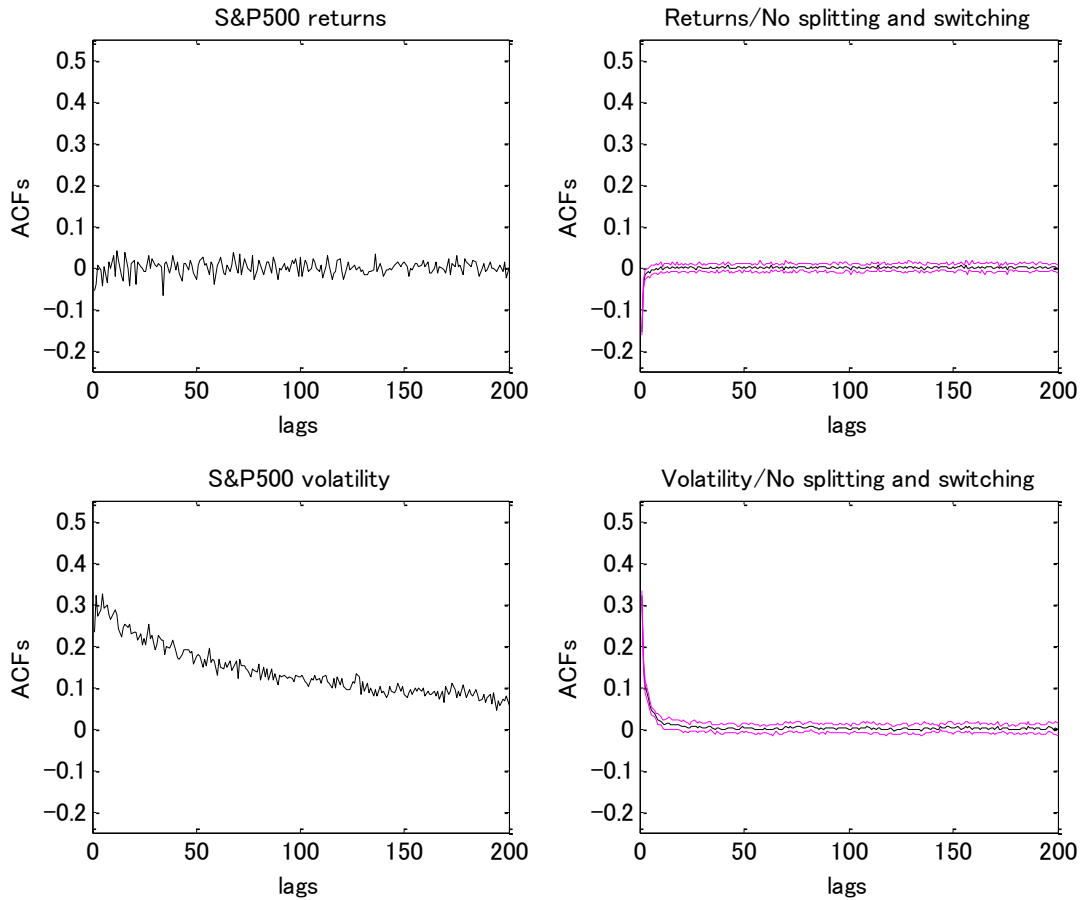


Figure 1: Autocorrelation functions (ACFs) of returns (top) and volatility (bottom) as a function of lags for the S&P500 (left) and in economies without order splitting and switching rules (right). Black lines in the benchmark economy represent the mean autocorrelation functions over 20 runs. Pink lines represent one standard deviations from the mean ACFs over 20 runs.

The returns of the S&P500 are negatively skewed in that downward price movements are of a greater magnitude than are the upward movements. We cannot conclude that our benchmark economy generates this feature, because the skewness is close to zero. The zero skewness comes from our setup (Equation (8)) where expectations are randomly determined and symmetric between upward and downward movements.

Figure 1 illustrates the autocorrelation functions of returns and volatility in our benchmark economy and S&P 500. The left top and bottom figures show that the actual stock market is informationally efficient such that the returns do not autocorrelate with different lag lengths; meanwhile, the volatility positively autocorrelates even with long lags, thus indicating evidence of clustered volatility. However, when the agents do not make use of dynamic switching and order-splitting strategies, we do not generate the slowly decaying autocorrelations of volatility, although the autocorrelations of returns still hover around zero.

We will show in the following subsections that both the order-splitting and dynamic switching strategies independently contribute to fatter tails and slower decaying autocorrelations of volatility than in our benchmark economy, and thus generate results that more closely approximate those in the real stock market.

3.1.2. Distributions of order patterns

Next, we introduce order patterns (absolute and then conditional frequency of order types) demonstrated in certain previous empirical studies. Table 3 summarizes the empirical findings of absolute frequency of order types in the Paris Bourse in Biais, Hillion, and Spatt (1995), the Toronto Stock Exchange in Griffiths, Smith, Turnbull, and White (2000), and the Australian Stock Exchange in Hall and Hautsch (2006) and Duong, Kalev, and Krishnamurti (2009). Each column presents the frequency of the sum of the buy and sell orders of each order type and constitutes a probability vector adding to 1. As described in the Introduction section, the empirical evidence of the absolute frequency of order submissions shows that market orders amount to approximately 28–53% of all orders submitted to the book, while limit orders within (outside) the spread amount to approximately 2–19% (27–65%).

Table 3: Absolute frequency of order types. The numbers are calculated from data on the Paris Bourse in Biais, Hillion, and Spatt (1995), the Toronto Stock Exchange in Griffiths, Smith, Turnbull, and White (2000), and the Australian Stock Exchange in Hall and Hautsch (2006) and Duong, Kalev, and Krishnamurti (2009).

	Biais et al. (1995)	Griffiths et al. (2000)	Duong et al. (2009)		Hall and Hautsch (2006)		
			Institutional orders	Individual orders	NCP	TLS	WOW
Market order	0.533	0.472	0.333	0.334	0.285	0.460	0.456
Limit order within spread	0.195	0.140	0.111	0.089	0.067	0.025	0.086
Limit order outside spread	0.271	0.388	0.556	0.577	0.648	0.515	0.458

Note: NCP, TLS, and WOW refer to News Corporation, Telstra, and Woolworths, respectively.

The empirical findings on the conditional frequencies of order types on the spread and depth are summarized as follows.

Property 1) Agents are likely to place less (more) aggressive orders as the bid–ask spread widens (narrows).

Property 2) Agents tend to submit more (less) aggressive orders as the *same side* of the book becomes thicker (thinner).

Property 3) Agents tend to submit less (more) aggressive orders as the *opposite side* of the book becomes thicker (thinner).¹²

Investors in real stock exchanges take the actions in line with the time and price priority rules in a limit order market: orders submitted earlier, buy orders at a higher limit price, and sell orders at a lower limit price are higher priorities for execution. As the depth becomes thicker, limit orders that are placed outside the spread are less likely to be executed as the risk of non

¹² Among several empirical studies, Property 2) is consistent with the empirical findings presented in each of Biais, Hillion, and Spatt (1995), Duong, Kalev, and Krishnamurti (2009), Griffiths, Smith, Turnbull, and White (2000), Hall and Hautsch (2006), Handa, Schwartz, and Tiwari (2003), and Rinaldo (2004). Property 3) has also been presented in the empirical works by Duong, Kalev, and Krishnamurti (2009), Griffiths, Smith, Turnbull, and White (2000), Hall and Hautsch (2006), Handa, Schwartz, and Tiwari (2003), and Rinaldo (2004).

execution increases. Thus, investors cease to place limit orders outside the spread and are more likely to submit market orders so as to acquire immediate executions. Investors reverse their strategies as the depth thins.

Table 4 summarizes our results regarding the absolute frequency of order types in our benchmark economy. The frequencies of placing market and limit orders are within the range found in the empirical evidence; however, this happens simply because in our setup, the limit price is symmetrically determined from the best prices. Therefore, the sum of the frequencies of the market orders and limit orders within the spread is almost equal to the frequency of the limit orders outside the spread. However, the following indicates that the conditional frequency of order types is not explained when we do not account for any strategic behavior within the modeling agents' strategies.

Table 4: Absolute frequency of order types: An economy without order-splitting and dynamic predictor selection ($N: 1,000$). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 25 runs.

Market order	Limit order within spread	Limit order outside spread
0.365 (0.003)	0.133 (0.003)	0.502 (0.003)

We extend the approach used by Biais, Hillion, and Spatt (1995) to investigate the statistical properties of the conditional frequency. The state of the order book is characterized by the spread and depth of the order book. The spread is defined as large (small) when it is larger (smaller) than the mean over the previous 25,000 clicks. A large (small) depth indicates that the depth is larger (smaller) than the mean over the previous 25,000 clicks.

Table 5 presents the results regarding the conditional frequency of order types in our benchmark economy. In this table, "Limit order outside spread" is the sum of the orders at the

best prices and below (above) the best bid (ask). Each row includes the frequency of the orders (adding to 1), conditional on the book's previous state. For example, the number in the first column in the first row of Table 5, Panel B indicates that the market order is observed by 18.2% when the depth on the same side declines. The number in the fourth column of the first row of Table 5, Panel B shows that the market order is observed by 18.1% when the depth on the opposite side is small. For this analysis, we use in each run the final 20,000 periods of the 3,000,000 clicks.

Table 5: Conditional frequency of order types given the state of the order book: An economy without order-splitting and dynamic predictor selection ($N: 1,000$). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs. Each row includes the frequency of the orders (adding to 1), conditional on the book's previous state.

Panel A:	Market order	Limit order within spread	Limit order outside spread
Small spread	0.429 (0.005)	0.069 (0.002)	0.502 (0.005)
Large spread	0.278 (0.004)	0.220 (0.006)	0.503 (0.005)

Panel B:	Same side			Opposite side		
	Market order	Limit order within spread	Limit order outside spread	Market order	Limit order within spread	Limit order outside spread
Small depth	0.182 (0.002)	0.067 (0.002)	0.252 (0.003)	0.181 (0.002)	0.067 (0.001)	0.251 (0.002)
Large depth	0.183 (0.002)	0.066 (0.002)	0.251 (0.002)	0.184 (0.002)	0.066 (0.002)	0.251 (0.003)

The results in Table 5 demonstrate that certain strategic behaviors may be required to replicate the conditional frequency of order types on depth. For example, the frequency of the market orders is 18.2% when the depth is small, but it remains virtually unchanged when the depth is large (18.3%).

In addition to the absolute and conditional frequencies of order types, the empirical distribution of the number of splits is another market microstructure feature that we need to reproduce in our order-splitting economy. Among the very limited number of studies that provide evidence, Vaglica, Lillo, Moro, and Mantegna (2008) study the Spanish stock exchange and find that the exponent for the number of splits is 1.8. In the following subsection on an order-splitting economy, we investigate what condition is required to generate the value of the scaling exponent, and how it is related to the clustered volatility and fat tails.

3.2. Simulation results in an order-splitting economy

We first demonstrate that an economy with an order-splitting strategy successfully generates clustered volatility and fat tails in an informationally efficient market. We then show that the number of splits, which are required to generate the return features, matches those found in empirical studies, such as that of Vaglica, Lillo, Moro, and Mantegna (2008). Furthermore, we demonstrate that the order-splitting economy is unable to reproduce the absolute and conditional frequencies of order types found in several previous empirical studies.

3.2.1. Empirical return features

In this subsection, we simulate an order-splitting economy in which all limit orders stored in the order book are observable to all traders. We vary the number of agents to between 250 and 2,500, from the original value of 1,000; we do this to investigate how the market size influences the return distribution. Table 4 shows the return distribution results, along with those from the S&P500 returns and returns from an economy without order-splitting and switching strategies. In addition, Figure 2 illustrates the average autocorrelation functions of returns and volatility over

20 runs as a function of lags. We highlight the following four results. First, although the return variations in our order-splitting model tend to be smaller than those from the S&P500, they are larger than those in our benchmark economy. In particular, we replicate the fat tail phenomenon in our order-splitting economy. For example, the kurtosis in our benchmark economy is 3.582, but it is 5.056 in our order-splitting economy with 1,000 agents. Second, in our benchmark economy, the autocorrelations of volatility collapse to zero for some short lags. However, the order-splitting strategy influences volatility dynamics, thereby generating highly persistent behavior in the series that resembles the characteristics of actual stock markets. This indicates that the order-splitting strategy may explain the clustered volatility in actual stock exchanges.

Table 6: Summary statistics of returns in economies with order splitting. The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs. N is the number of agents in the market.

	Mean $\times 10000$	Std. $\times 100$	Skewness	Kurtosis
S&P500	2.719	1.134	-0.240	11.81
No order splitting and switching	0.005 (0.022)	0.048 (0.001)	-0.002 (0.028)	3.582 (0.054)
Order splitting ($N: 250$)	-0.003 (0.058)	0.056 (0.004)	-0.006 (0.039)	4.352 (0.659)
Order splitting ($N: 1,000$)	0.007 (0.057)	0.063 (0.009)	0.010 (0.029)	5.056 (1.119)
Order splitting ($N: 2,500$)	0.015 (0.067)	0.068 (0.012)	0.008 (0.044)	5.157 (1.253)

Third, as in our benchmark economy, the autocorrelations of returns hover around zero even with very short lags, indicating that the markets are informationally efficient. This result derives from the influence of the bid–ask bounce: although the prices in our order-splitting economies can be persistent, once a market order enters the opposite side of the order book, the price returns to the best price on the opposite side, thus restoring market efficiency.

Fourth, our results on return features are robust with different order-splitting market sizes. However, as shown in Figure 2, we observe a stronger persistence of volatility when the market size becomes large (i.e., $N = 2,500$). This result indicates that in terms of generating clustered volatility, the order-splitting strategy is more crucial in large capital stocks.




Figure 2: Autocorrelation functions (ACFs) of returns (top) and volatility (bottom) as a function of lags in economies with order splitting with different market sizes

The following are a few possible interpretations of our results. Our order-splitting mechanism creates a persistence in order flows from a given agent with a small likelihood of

other agents placing new orders. This persistence is not observed in an economy without the order splitting due to the higher probability of a new agent entering the market sooner and placing his or her market order of the opposite sign or limit orders, thereby tending to collapse the persistence of the order signs. In our order-splitting economy, when an agent enters the market with a large order, he or she persistently executes the same sign of market orders that generates a persistence of order signs. When agents continue to do so, they persistently remove liquidity from the order book, so that the order-book depth tends to thin. A thinner order book generates larger price changes, and a persistently thinner order book will bring about persistently larger price changes. However, when an agent submits a large limit order, the order book becomes thicker, and as a result, price changes shrink, and these smaller price changes persist. Therefore, we can generate larger price changes, followed by larger price changes and smaller price changes, which are later followed by smaller price changes. As the limit orders are removed from the book and the best prices change, the change in the best prices sometimes becomes large, as there are no limit orders at limit prices next to the previous best prices. This generates unusual price changes that comprise the origin of the fat tails. Such persistent behavior in order flows does not emerge in an economy without order-splitting strategy, because agents do not split their orders and thus execute their orders at once so that new agents frequently enter the market. Therefore, we do not observe clustered volatility and fat tails when agents do not adopt an order-splitting strategy.

Our results suggest that when the number of agents in the market is large, it is likely that the depth of the order book will become large. As a result, when market orders persistently enter the order book, they are repeatedly executed at the same best prices. For this reason, the best prices do not frequently or markedly change, even when market orders do change the best prices.

Thus, we are more likely to observe slower but persistent price changes in an order-splitting economy with a larger market size. Therefore, the autocorrelations decay more slowly in our order-splitting economy when the number of agents grows. The higher probability of less-frequent price changes means a smaller occurrence of certain large price changes, which in turn contributes to producing fatter tails in the return distribution of an economy that features a larger number of agents. Therefore, our results indicate that an order-splitting strategy contributes more to clustered volatility and fat tails in larger capital stocks.

Subsection 3.2.2 examines the distributions of the number of splits, and which distributions are necessary to simultaneously generate clustered volatility and fat tails; it also compares those distributions to those found in real stock exchanges.

3.2.2. Distributions of the splits

This subsection demonstrates that our order-splitting model can generate the distribution of the number of splits that matches that shown in the empirical evidence from real stock markets. We also show that certain single stock traders need to split their orders into very large pieces, in order to reproduce clustered volatility and fat tails.

Assuming that the number of splits is power-law distributed, the probabilities that the number of splits n^i is larger than x are given by:

$$P(n^i > x) \sim x^{-\zeta}, \quad (18)$$

where ζ is the scaling exponent for the split number distribution. We estimate ζ in each of the number of traders (250, 1,000, and 2,500). The results are presented in Table 7; this table also shows the means of the number of splits and fractions of splits of more than 10, 100, and 1,000.

Table 7: Summary statistics of the number of splits and the total order size: Order-splitting market. The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs.

Number of agents: N	Exponents for the number of splits: ζ	Mean	% of splits of more than		
			10	100	1,000
250	1.60 (0.16)	249.4	0.57	0.33	0.06
1,000	1.60 (0.43)	229.3	0.57	0.33	0.05
2,500	1.84 (0.22)	177.5	0.55	0.31	0.04

The results are summarized as follows. First, the number of splits tends to be smaller, on average, as the market size grows. For example, in an economy with 250 agents, the mean of the splits is 249.4, while the percentages of the splits of more than 10, 100, and 1,000 are 57%, 33%, and 6%, respectively. However, when the number of agents increases to 2,500, the mean of the splits amounts to 177.5, and the percentages of the splits of more than 10 and 1,000 become 55% and 4%, respectively. The changes in the results lead to an increase in the scaling exponents of the number of splits. The scaling exponent is 1.60 in an economy with 250 agents, while it is 1.84 when the number of agents grows to 2,500. The change is attributed to our setup, whereupon an agent who starts to split his or her order has a greater chance of placing his or her subsequent splits as the total amount of orders waiting to be split diminishes. The chances are likely to be higher as the number of agents increases as indicated in Equations (11) and (12). On the other hand, as demonstrated in the previous subsection, clustered volatility and fat tails are replicated in all economies with 250, 1,000, and 2,500 agents.

Second, our scaling exponents of the number of splits necessary to generate clustered volatility and fat tails match that presented in the empirical evidence by Vaglica, Lillo, Moro, and Mantegna (2008). In particular, our estimate of the number of splits is 1.84 when there are 2,500 agents in the market, while that by Vaglica, Lillo, Moro, and Mantegna (2008) is 1.8. Our

agents endogenously determine the size and thus their number of their splits, in order to minimize market impact risk. Therefore, our result suggests that we have endogenously replicated the empirical feature on the number of splits and thus provided a theoretical explanation as to why real investors tend to have such a large number of splits.

3.2.3. Absolute and conditional frequencies of order types

This subsection investigates the absolute frequency of order submissions and the conditional frequency of order types on the depth and spread in our order-splitting economy that replicates the clustered volatility, fat tails, and empirical scaling exponent of the number of splits in the previous section. This investigation is carried out by comparing the absolute and conditional frequencies in our order-splitting economy with those in real stock exchanges.¹³ We demonstrate that the order-splitting economy cannot replicate the absolute and conditional frequencies of the order types observed in real stock exchanges, and thus provide evidence that clustered volatility and fat tails originate not only from the order-splitting strategy, but also from other strategies in actual stock exchanges.

Table 8 summarizes our results regarding the absolute frequency in our order-splitting economy with 1,000 agents. Compared to empirical evidence regarding the absolute frequency of order types (Table 3), market orders are placed much more frequently (99.7%) and limit orders within and outside the spread are submitted much less frequently (0.1% and 0.3%, respectively), compared to in real stock exchanges.¹⁴

¹³ We provide only the results with 1,000 agents, as the order pattern results are quite robust with different numbers of agents.

¹⁴ Our results indicate that each split size is sufficiently small so as not to deplete all limit orders in the order book even when market orders are executed with limit orders in the book at a frequency of 99.7%.

Table 8: Absolute frequency of order types: Order-splitting economy (N: 1,000). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs.

Market order	Limit order within spread	Limit order outside spread
0.997 (0.002)	0.001 (0.0005)	0.003 (0.001)

Table 9: Conditional frequency of order types given the state of the order book: Order-splitting economy (N: 1,000). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs. Each row includes the frequency of the orders (adding to 1), conditional on the book's previous state.

Panel A:	Market order	Limit order within spread	Limit order outside spread
Small spread	0.997 (0.002)	0.0007 (0.0004)	0.002 (0.001)
Large spread	0.996 (0.002)	0.001 (0.0006)	0.003 (0.001)

Panel B:	Same side			Opposite side		
	Market order	Limit order within spread	Limit order outside spread	Market order	Limit order within spread	Limit order outside spread
Small depth	0.493 (0.052)	0.0006 (0.0003)	0.001 (0.0006)	0.503 (0.051)	0.0005 (0.0002)	0.001 (0.0006)
Large depth	0.510 (0.045)	0.0004 (0.0003)	0.001 (0.0005)	0.487 (0.045)	0.0005 (0.0003)	0.001 (0.0006)

Table 9 presents our result regarding the conditional frequency of order types. As shown in Table 9, none of the properties of the conditional frequencies are replicated in our order-splitting economy. For example, in Table 9, Panel A, market orders are frequently placed regardless of the size of the spread, and the market order frequencies are similar (99%) when the spread declines and expands. In Table 9, Panel B, when the depth on the same side becomes thick, market orders enter the order book at a frequency (51%) similar to that under a thinner depth (49.3%). Even when the depth on the same side becomes thicker, the limit order placements do not increase. In addition, order placement frequency is very similar when the opposite side of the order book has a larger or smaller depth. For example, the frequency of limit

orders inside the spread is 0.05%, regardless of the degree to which the depth on the opposite side of the order book changes. In our order-splitting economy, our agents persistently execute small splits so as to mitigate market impact risk. This action is observed even when the spread increases or the depth changes.

Our results in Subsection 3.2.2 show that an order-splitting strategy, which is consistent with the empirical distribution on the number of splits, can explain clustered volatility and fat tails. However, our results on the order patterns contradict the empirical facts, suggesting that other trading strategies may need to be considered if we are to explain all of the empirical features in reality. Subsection 3.3 examines another popularly utilized strategy namely, a dynamic switching strategy, and evaluates whether the switching strategy has greater explanatory power regarding the empirical features of both of the returns and market microstructure.

3.3. Simulation results on dynamic predictor selection

We have seen in Subsection 3.1 that our benchmark economy, which is the one without the order splitting and dynamic predictor selection, does not generate clustered volatility and fat tails. Agents' expectations in the economy depend on the noise-trading component. We relax the assumption, and the expectations in this subsection are formed by dynamic predictor selection between fundamental and trend-following strategies, as described in Section 2. We analyze the influence of the switching strategy exclusively of the effect from the order splitting, and confirm that the model does produce clustered volatility and fat tails, even in the absence of the order-splitting mechanism. However, we show that the empirical properties of conditional frequency of order types do not resemble those found in previous empirical studies. Thus, we conclude that

while the switching strategy does help explain clustered volatility and fat tails, it is not the sole cause of all empirical phenomena in the financial markets.

3.3.1. Empirical return features

Table 10 presents summary statistics of returns in an economy with dynamic predictor selection when the number of agents N equals 250, 1,000, and 2,500. We also list there the return properties from real S&P500 data, as well as from our simulated market without order-splitting and switching strategies (Table 2), so as to highlight the impact of the dynamic predictor selection strategy. Figure 3 illustrates the persistence of returns and volatility in the switching economy. The information in Table 10 and Figure 3 confirm that the dynamic predictor selection strategy does produce clustered volatility and fat tails in an informationally efficient market, although the return variations are lower in the switching economy than in the actual market. The standard deviations of the returns in the switching and actual markets are 0.057/100-0.058/100 and 1.134/100, respectively. However, in an economy with the dynamic predictor selection, the kurtosis is 4.545-5.163. Additionally, the return volatility is auto-correlated with long lags.

Table 10: Summary statistics of returns in economies with dynamic predictor selection. The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs.

	Mean \times 10000	Std. \times 100	Skewness	Kurtosis
S&P500	2.719	1.134	-0.240	11.81
No order splitting and switching	0.005 (0.022)	0.048 (0.001)	-0.002 (0.028)	3.582 (0.054)
Switching strategy (N: 250)	0.004 (0.033)	0.058 (0.004)	0.003 (0.044)	5.163 (0.909)
Switching strategy (N: 1,000)	0.006 (0.022)	0.057 (0.003)	-0.001 (0.035)	4.765 (0.466)
Switching strategy (N: 2,500)	-0.005 (0.029)	0.057 (0.003)	0.015 (0.015)	4.545 (0.379)

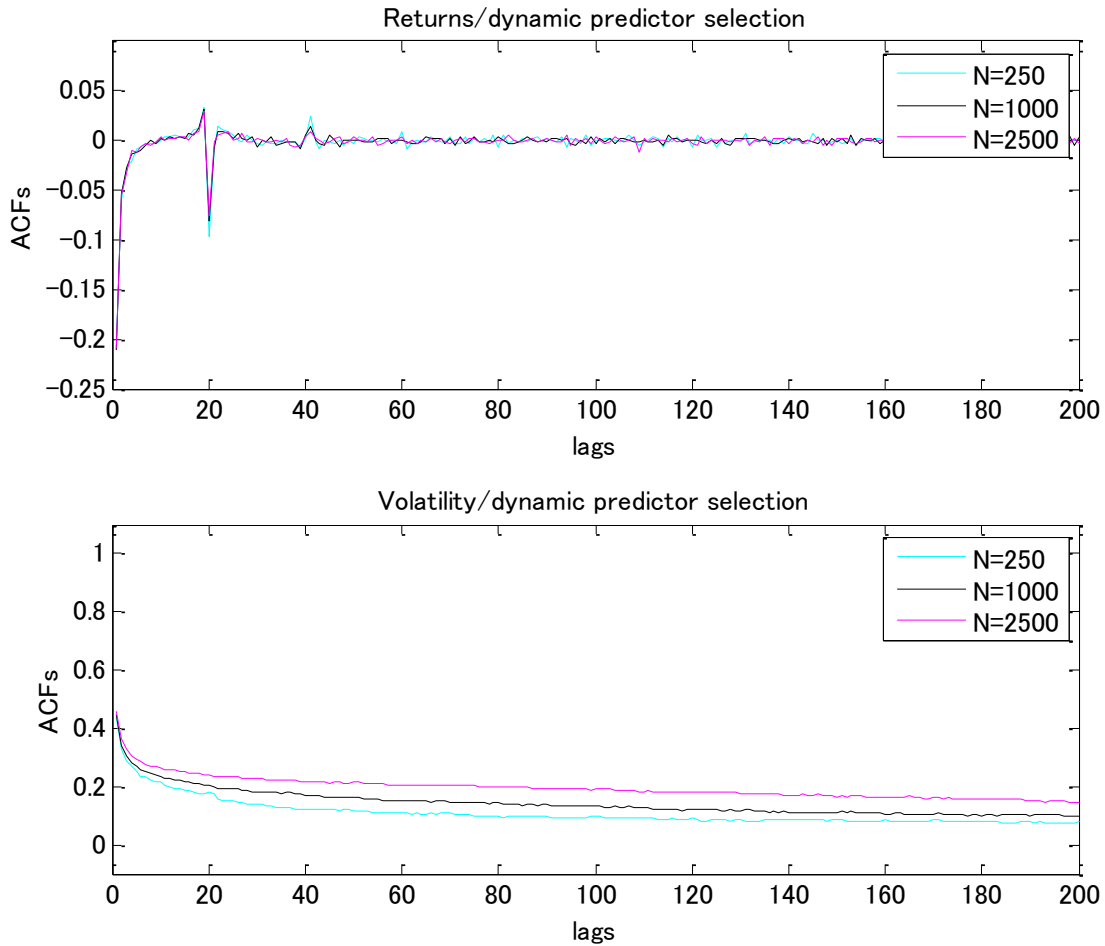


Figure 3: Autocorrelation functions (ACFs) of returns (top) and volatility (bottom) as a function of lags in an economy with dynamic predictor selection where $N = 250, 1,000,$ and $2,500$

As explained in the previous agent-based model, when most agents tend to choose a trend-following rule, the price follows the past trend, meaning that the price increases when they buy and decreases when they sell. Then, the forecast errors from the trend-following strategy decline and thus more agents are likely to select the trend-following rule. As a result, the price trend intensifies, exacerbating volatility. This higher-volatility trend is persistent, as agents tend to adhere to the same rule over a protracted period. Our result indicates that this intensified price

movement is the cause of fat tails in the return distribution. As the price deviates markedly from the fundamental price, a certain number of agents predict that the price will revert to the fundamental price. The action changes the price patterns so that they no longer follow past trends, and volatility starts to stabilize. This price movement increases the profit stemming from the fundamental predictor, and more agents start to select the fundamental strategy; this further reduces volatility, and smaller price movements become persistent. The price thus starts to move back to the fundamental price, and market efficiency is achieved. The returns are negatively auto-correlated with a few lags, on account of the bid-ask bounce phenomenon. In our limit order market, the agents buy at the best ask, but they sell at the best bid, meaning that the transaction prices bounce back and forth as they trade, thus generating a negative persistence in transaction price.

Our dynamic predictor selection results are robust in terms of economy size. The autocorrelations of returns are very similar but those of volatility become more persistent when there are 2,500 agents in the market. On the one hand, as more agents switch so as to follow trend-following rules, the price trend intensifies. On the other hand, price fluctuations tend to be smaller, as the majority of agents tend to use the fundamental strategy. The magnitudes of volatility change on account of the switching behavior of agents. However, our results suggest that changes in the magnitudes of volatility are likely to occur more slowly when the number of agents is larger, and each agent action thus has less of an influence on the macroeconomic variables. The result indicates that dynamic predictor selection is the cause of clustered volatility and fat tails in an informationally efficient market, and that it is more crucial in larger capital stocks.

3.3.2. Distributions of order patterns

The information in Tables 11 and 12 demonstrates that an economy with dynamic predictor selection produces an absolute frequency of order choice; this is similar to the frequency in a real economy, but the conditional frequency of order types does not match that seen in real-world data. For example, market orders are submitted to the order book with 39.1% of the total orders placed, while limit orders outside the spread account for 50.2% of all orders. The fractions are quite consistent with those observed in past empirical studies on order placement. In addition, our agents are more likely to submit market orders when the spread becomes small: the percentage of market order placement is 29.4% when the spread is large, but that figure increases to 44.9% when the spread narrows.

Table 11: Absolute frequency of order types: Dynamic predictor selection (N: 1,000). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 25 runs.

Market order	Limit order within spread	Limit order outside spread
0.391 (0.027)	0.107 (0.028)	0.502 (0.003)

Table 12: Conditional frequency of order types given the state of the order book: Dynamic predictor selection (N: 1,000). The bold numbers in our economies are means over 20 runs, while the numbers in parentheses are standard deviations over 20 runs. Each row includes the frequency of the orders (adding to 1), conditional on the book's previous state.

Panel A:	Market order	Limit order within spread	Limit order outside spread
Small spread	0.449 (0.014)	0.050 (0.014)	0.502 (0.005)
Large spread	0.294 (0.053)	0.205 (0.054)	0.502 (0.004)

Panel B:	Same side			Opposite side		
	Market order	Limit order within spread	Limit order outside spread	Market order	Limit order within spread	Limit order outside spread
Small depth	0.200 (0.020)	0.053 (0.013)	0.257 (0.011)	0.191 (0.011)	0.052 (0.016)	0.246 (0.010)
Large depth	0.188 (0.013)	0.054 (0.016)	0.244 (0.012)	0.200 (0.021)	0.056 (0.013)	0.258 (0.014)

However, the conditional frequency on the order-book depth is not consistent with that seen in the actual market. In particular, the conditional frequency does not change when the same side and the opposite side of the order book become thicker. The frequencies of market order placement are 20% for a large depth of the same side, and 18.8% for a small depth. 19.1% of the orders placed are market orders when the opposite side of the order book is small, but that figure remains virtually unchanged (20%) when the depth on the opposite side becomes large. The model with dynamic predictor selection does not replicate the conditional frequency, because agents in the model form their expectations based only on price information not on the order-book condition.

We have demonstrated that dynamic predictor selection and order splitting can independently be a cause of clustered volatility and fat tails in an informationally efficient market. However, neither of strategy types reproduces the empirical features of order choice. The results indicate that in addition to dynamic predictor selection and order splitting, certain other

strategies may concurrently operate in reality to reproduce the market microstructure feature on the order types. Our order-splitting agents slice their orders into small pieces by mitigating market impact risk. With our dynamic predictor selection, the agents look only at past price information. Our results suggest that market impact risk is not the only risk—and past price information not the only information—that actual investors examine when forming their strategies. For example, there is a certain strategy wherein investors base their strategies on the demand–supply condition in the order book and take into account non-execution and/or picking-off risks (Yamamoto, 2011); this is a possible strategy that real investors could actually use, and it stands as a possible candidate for explaining empirical phenomena on the order patterns.

4. Conclusion

In this study, we have demonstrated that order-splitting and dynamic switching strategies successfully can explain clustered volatility and fat tails of return distribution in an informationally efficient market. However, we have also confirmed that dynamic switching strategies do not replicate conditional frequencies of order submissions, while the order-splitting strategy cannot sufficiently explain both the absolute and conditional frequencies of order submissions. Therefore, our theoretical results suggest that order splitting and dynamic predictor selection operate simultaneously in reality, but may not be the only strategies at work in the market. We also find that the persistence of volatility increases with a greater number of agents in both types of markets, implying that order splitting and dynamic predictor selection are more crucial for the larger capital stocks in real stock markets. Our results imply that in addition to the order-splitting or dynamic switching strategy, other strategies—such as that constructed on the basis of the order-book information and determined while bearing in mind non-execution and

picking-off risks—may be the main generators of both macroeconomic and market microstructure properties.

Furthermore, identifying empirically which strategies people tend to choose, as well as the economic conditions in which they use them can serve as the subjects of our future research. As a possible first step in addressing these topics, it would be necessary to collect survey data, or data from laboratory experiments that make use of human subjects including individual traders' identities which track who make use of order-splitting, dynamic predictor selection, or any other strategy, such as any that the order-book information. In addition, it will also be important to conduct more theoretical research that discriminates and evaluates behavioral assumptions other than order splitting and dynamic predictor selection within a theoretical model for explaining the empirical features. All of these are important research paths in the area of agent-based modeling.

References

- Alfarano, S., Lux, T., Wagner, F., 2008. Time variation of higher moments in financial markets with heterogeneous agents: An analytical approach. *Journal of Economic Dynamics and Control* 32, 101-136.
- Alfonsi, A., A. Fruth, and A. Schied, 2010. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance* 10, 143–157.
- Almgren, R. and N. Chriss, 2000. Optimal execution of portfolio transactions. *Journal of Risk*. 3, 5–39.
- Almgren, R. and J. Lorenz, 2007. Adaptive arrival price. In *Algorithmic Trading III: Precision, Control, Execution*, edited by B.R. Bruce, (Institutional Investor Journals: New York).
- Anufriev M., Arifovic J., Ledyard J., Panchenko V., 2012. Efficiency of continuous double auctions under individual evolutionary learning with full or limited information, *Journal of Evolutionary Economics*, forthcoming.
- Anufriev, M., and V. Panchenko 2009. Asset Prices, Traders' Behavior and Market Design," *Journal of Economic Dynamics and Control* 33, 1073-1090.
- Bertsimas, D., A. Lo, 1998. Optimal control of execution costs. *Journal of Financial Markets* 1, 1–50.
- Bessembinder, H., Panayides, M.A., Venkataraman, K., 2009. Hidden liquidity: An analysis of order exposure strategies in electronic stock markets. *Journal of Financial Economics* 94, 361–383.
- Biais, B., Hillion, P., Spatt, C., 1995. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance* 50, 1655–1689.
- Bouchaud, J.-P., Gefen, Y., Potters, M., Wyart, M., 2004. Fluctuations and response in financial markets: The subtle nature of 'random' price changes. *Quantitative Finance* 4, 176-190.
- Brock, W.A., Hommes, C.H., 1997. A rational route to randomness. *Econometrica* 65, 1059–1095.
- Brock, W., Hommes, C., 1998. Heterogeneous beliefs and routes to chaos in a simple asset pricing model. *Journal of Economic Dynamics and Control* 22, 1235–1274.
- Chiarella, C., He, X-Z., Wei, L., 2015. Learning, information processing and order submission in limit order markets. Working paper, University of Technology at Sydney.
- Chiarella, C., Iori, G., Perelló, J., 2009. The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Economic Dynamics and Control* 33, 525-537.

- Chen, S.-H., Chang, C.-L., Du, Y.-R., 2012. Agent-based economic models and econometrics. *Knowledge Engineering Review* 27(2), 187–219.
- Ding, Z., Granger, C., Engle, R., 1993. A long memory property of stock market returns and a new model. *Journal of Empirical Finance* 1, 83-106.
- Duong, H. N., Kalev, P., Krishnamurti, C., 2009. Order aggressiveness of institutional and individual investors. *Pacific-Basin Finance Journal* 17, 533-546.
- Farmer, J.D., Lillo, F., 2004. On the origin of power laws in financial markets. *Quantitative Finance* 4, C7–C10.
- Foucault, T., O. Kadan, and E. Kandel, 2005. Limit order book as a market for liquidity, *Review of Financial Studies* 18, 1171-1217
- Frankel, J.A., Froot, K.A., 1990. Chartists, fundamentalists and the demand for dollars, in: Courakis, A.S., Taylor, M.P. (eds.), *Private Behaviour and Government Policy in Interdependent Economies*. Oxford University Press, New York, pp. 73-126.
- Griffiths, M., Smith, B., Turnbull, A., White, R., 2000. The costs and the determinants of order aggressiveness. *Journal of Financial Economics* 56, 65–88.
- Goettler, R., C. Parlour, U. Rajan, 2005, Equilibrium in a dynamic limit order market, *Journal of Finance* 60, 2149–2192.
- Gould, M., Porter, M., Williams, S., Fenn, D., Howison, S. D., 2013. Limit order books. *Quantitative Finance* 13(11), 1709–1742.
- Hall, A., Hautsch, N., 2006. Order aggressiveness and order book dynamics. *Empirical Economics* 30, 973-1005.
- Handa, P., Schwartz, R., Tiwari, A., 2003. Quote setting and price formation in an order driven market. *Journal of Financial Markets* 6, 461–489.
- Hommes, C.H., 2006. Heterogeneous agent models in economics and finance, in: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics*, vol. 2: Agent-Based Computational Economics. North-Holland, Amsterdam, pp. 1109–1186.
- Kirman, A., 1993. Ants, rationality, and recruitment. *Quarterly Journal of Economics* 108, 137-156.
- LeBaron, B., 2006. Agent-based computational finance. in: Tesfatsion, L., Judd, K.L. (Eds.), *Handbook of Computational Economics*, vol. 2: Agent-Based Computational Economics. North-Holland, Amsterdam, pp. 1187–1233.

- LeBaron, B., Arthur, B., Palmer, R., 1999. Time series properties of an artificial stock market. *Journal of Economic Dynamics and Control* 23, 1487-1516.
- Lillo, F., Farmer, D., 2004. The long memory of the efficient market. *Studies in Nonlinear Dynamics & Econometrics* 8, Article 1.
- Lillo, F., Mike, S., Farmer, D., 2005. Theory for long memory in supply and demand. *Physical Review E* 71, 066122.
- Lorenz, J. and R. Almgren, 2011. Mean-variance optimal adaptive execution. *Applied Mathematical Finance* 18, 395-422.
- Lui, Y-H., Mole, D., 1998. The use of fundamental and technical analysis by foreign exchange dealers: Hong Kong evidence. *Journal of International Money and Finance* 17, 535-545.
- Lux, T., 1995. Herd behavior, bubbles and crashes. *The Economic Journal* 105, 881-896.
- Lux, T., 2001. The limiting extremal behaviour of speculative returns: an analysis of intra-daily data from the Frankfurt stock exchange. *Applied Financial Economics* 11, 299-315.
- Mandelbrot, B., 1972. Statistical methodology for non-periodic cycles: from the covariance to R/S analysis. *Annals of Economic and Social Measurements* 1, 259-290.
- Menkhoff, L., 2010. The use of technical analysis by fund managers: international evidence, *Journal of Banking and Finance* 34, 2573-2586.
- Menkhoff L., Taylor, M., 2007. The obstinate passion of foreign exchange professionals: Technical analysis. *Journal of Economic Literature* 45, 936-972.
- Mike, S., Farmer, D., 2008. An empirical behavioral model of liquidity and volatility. *Journal of Economic Dynamics and Control* 32, 200-234.
- Obizhaeva, A. and J. Wang, 2005. Optimal trading strategy and supply/demand dynamics. NBER Working Paper No. 11444.
- Parlour, C., 1998. Price dynamics in limit order markets. *Review of Financial Studies* 11, 789-816.
- Perold, A. F., 1988. The implementation shortfall: Paper versus reality. *Journal of Portfolio Management* 14, 4-9.
- Rinaldo, A., 2004. Order aggressiveness in limit order book markets. *Journal of Financial Markets* 7, 53-74.
- Roşu, I., 2009. A dynamic model of the limit order book. *Review of Financial Studies* 22, 4601-4641.

Sussman, A. 2005. Institutional Equity Trading 2005: a Buy-Side Perspective. *Tabb Group Research Report*.

Tóth, B., I. Palit, F. Lillo, and D. Farmer. 2015. Why is order flow so persistent?, *Journal of Economic Dynamics and Control* 51, 218–239.

Vaglica, G., Lillo, F., Moro, E., Mantegna, R., 2008. Scaling laws of strategic behaviour and size heterogeneity in agent dynamics. *Physical Review E* 77, 036110.

Yamamoto, R., LeBaron, B., 2010. Order-splitting and long-memory in an order-driven market. *European Physical Journal B* 73, 51-57.

Yamamoto, R., 2011. Order aggressiveness, pre-trade transparency, and long memory in an order-driven market. *Journal of Economic Dynamics and Control* 35, 1938-1963.