Available online at www.sciencedirect.com
ScienceDirect
Physica A 383 (2007) 85–89
PHYSICA A
www.elsevier.com/locate/physa

Long-memory in an order-driven market

Blake LeBaron[a], Ryuichi Yamamoto[b,*]

[a]International Business School, Brandeis University, Waltham, MA 02453, USA
[b]Department of International Business, National Chengchi University, Taipei 116, Taiwan

Available online 1 May 2007

Abstract

This paper introduces an order-driven market with heterogeneous investors, who submit limit or market orders according to their own trading rules. The trading rules are repeatedly updated via simple learning and adaptation of the investors. We analyze markets with and without learning and adaptation. The simulation results show that our model with learning and adaptation successfully replicates long-memories in trading volume, stock return volatility, and signs of market orders in an informationally efficient market. We also discuss why evolutionary dynamics are important in generating these features.

Keywords: Microstructure; Agent-based; Long-memory; Order flow

1. Introduction

This paper provides a simulation analysis of an order-driven market to explain the following empirical properties. First, trading volume is persistent over time [1]. Second, stock return volatility has long-memory, meaning that the stock price fluctuations have positive autocorrelations (e.g., [2]). Third, there is a positive contemporaneous correlation between volatility and volume [3]. Finally, the signs of market orders follow a long-memory process; yet the market is informationally efficient [4] in that returns are uncorrelated over time.

We modify an order-driven market, which was originally constructed by Chiarella and Iori [5]. In their order-driven market, heterogeneous investors submit limit or market orders according to exogenously fixed rules. Our key modification is that the trading rules are repeatedly updated via learning and adaptation of the investors. At a given point in time, investors look back at the past performance of their predictions. They realize their past mistakes on these predictions and update their prediction methods to new ones by *imitating* strategies of more successful investors. The success of others is measured by their past predictive accuracy. Thus, investors learn from the past and adapt their behavior to improve their future performances in the market. We compare the dynamics that emerge from this economy with and without imitation. Our economy with certain types of evolutionary behavior, such as imitation, successfully replicates the previously mentioned four empirical regularities on the stock market. We will present evidence that suggests the imitation of strategies is key to understanding these empirical properties in an order-driven market.

*Corresponding author.
E-mail address: ryuichi@nccu.edu.tw (R. Yamamoto).

0378-4371/$ - see front matter © 2007 Elsevier B.V. All rights reserved.
doi:10.1016/j.physa.2007.04.090

Several recent papers have attempted to replicate features of an order-driven market with an environment with zero intelligence agents (e.g., [6]). In these models, random order flows are sent to the market clearing mechanism, and the resulting time series are analyzed. These random behavior models are an important benchmark for understanding which time series features could be driven by institutional features alone. We move from the zero intelligence world to one of "low intelligence", or bounded rationality for several reasons. First, we are interested in having some economic motivation for our agents generating buy and sell orders. This could protect us from results being driven by wildly irrational trading behavior. Second, we seek an endogenous explanation for the extreme persistence in order signs, volume, and liquidity. Finally, our agents are motivated by economic objectives to at least work hard to ensure an efficient market. We are interested in seeing how efficient these markets can be in terms of prices, while generating the other persistent time series results. We will do this in a single market with simple, but adaptively rational agents.

The rest of the paper proceeds as follows. Section 2 presents the market structure. Section 3 gives simulation analyses and the results. The last section concludes.

2. Market structure

This section describes an artificial stock market based on the market outlined in Chiarella and Iori [5]. We examine a continuous double auction market where agents submit orders on the market sequentially. Orders are matched and executed according to the time and price priorities. A single asset is traded in the market, and its fundamental value is assumed to be constant. Agents are assumed to know the fundamental value p^f and the past history of the prices. The price p_t is determined at the level where any transaction occurs. When there is no transaction, a proxy of the price is given by the average of the lowest ask (the quoted ask, a_t^q) and the highest bid (the quoted bid, b_t^q). We use the number of events as a measure of time.

Agents observe all the past price series and determine their strategies based on their expectations of future prices. They make their predictions based on the following equation:

$$\hat{r}^i_{t,t+\tau} = g^i_1\left(\frac{p^f - p_t}{p_t}\right) + g^i_2\bar{r}_{L_i} + n_i\varepsilon_t, \tag{1}$$

where g_1^i, g_2^i, and n_i are fundamentalist, chartist, and noise-induced components, respectively, and initially these are randomly assigned according to the distributions; $g_1^i \sim |N(0,\sigma_1)|$, $g_2^i \sim N(0,\sigma_2), n_i \sim |N(0,n_0)|$, $\varepsilon \sim N(0,1)$. The value, $\bar{r}_{L_i}$, is given as

$$\bar{r}_{L_i} = \frac{1}{L_i}\sum_{j=1}^{L_i}\frac{p_{t-j} - p_{t-j-1}}{p_{t-j-1}}. \tag{2}$$

It is the average return over the interval, L_i, which is randomly drawn from a uniform and independent distribution across agents over the interval (1, $L_{\max}$).

Agents expect future prices at time $t+\tau$ according to those of the future returns as

$$\hat{p}^i_{t+\tau} = p_t e^{\hat{r}^i_{t,t+\tau}}. \tag{3}$$

Agents update the current price in their forecasting model at every click. An agent buys (sells) one unit of stock at a price $b_t^i(a_t^i)$ when she expects an increase (decrease) in the price. b_t^i and a_t^i are given as follows

$$b^i_t = \hat{p}^i_{t,t+\tau}(1 - k^i)$$
$$a^i_t = \hat{p}^i_{t,t+\tau}(1 + k^i), \tag{4}$$

where k^i is randomly and uniformly assigned across agents in the interval (0, $k_{\max}$), with $k_{\max} \leqslant 1$. k^i is fixed over time, but varies over agents. Agents are assumed to submit a limit buy (sell) order at $b_t^i(a_t^i)$ for one unit, if $b_t^i(a_t^i)$ is smaller (larger) than the current quoted ask, a_t^q (bid, b_t^q). If the bid (ask) price is higher (lower) than the best ask (bid) then the order executes immediately at the current best-quoted price on the book, and its price is recorded. The order is removed from the book once executed or expired at its lifetime τ. Agents make trading decisions sequentially. One trading round ends, once all have had a chance to submit their orders. At the end of each trading round, we scramble the ordering of agents' arrivals. Then, another trading round starts. After