An Agent-Based Model of Excess Kurtosis in Financial Market Returns

Neil Rayner[1], Steve Phelps[2], and Nick Constantinou[1]

[1] Essex Business School (EBS), University of Essex,
Colchester, CO4 3SQ, UK.
njwray@essex.ac.uk
[2] Centre for Computational Finance and Economic Agents (CCFEA),University of Essex,
Colchester, CO4 3SQ, UK.

Abstract. It is well known that the changes in the logarithmic price (returns) in financial time series data do not follow a Gaussian distribution when sampled at sufficiently high frequency. Researchers have conjectured that the excess kurtosis observed in actual high-frequency returns is caused by trend following behaviours; traders imitate each other, forming bubbles where prices diverge from a fundamental understanding of an asset's underlying value. We test this conjecture with an agent-based model of the financial market and find that excess kurtosis is directly related to the relative weighting of fundamentalist, trend following and noise trading strategies. All three strategies have an impact on kurtosis with the fundamentalist and trend following strategies increasing kurtosis, while noise trading reduces it.

Keywords: Return kurtosis, Agent-based models, Stylized facts

1 Introduction

Agent-based models are able to capture complex trading behaviour and market micro-structure which is difficult to incorporate in traditional financial equation-based models. More realistic models of trading agent decision making can be constructed, using heuristics which are consistent with empirical observations from the controlled study of actual human subjects ([5]).

One approach to validating agent-based models is to demonstrate that they produce simulated time series data which are consistent with the empirically-observed stylized facts of actual financial markets, and that these characteristics are insensitive to the settings of model free parameters; that is, we attempt to show that the model is robust. Given a robust agent-based model we can suggest real world mechanisms that may be responsible for empirically observed phenomena. We focus on a well known stylized fact of high-frequency time series data observed in real financial markets and we analyse to what extent different model assumptions are consistent with this phenomena. In this research, we test the hypothesis that the stylized fact of high kurtosis for return distributions is as a result of forecasting strategies employed by trading agents.

2 The Model

We adopt an existing model that has been the basis of much research in the agent-based modelling community ([6][7][12][3][11])). It is a non-adaptive expectations model with three classes of strategy which are used to form expectations about future returns (Chiarella and Iori [2]):

1. *fundamentalists* value a stock through an understanding of its hypothetical underlying value, in other words, based on expectations of the long term profitability of the issuing company;
2. *chartists* form valuations inductively from historical price data; and
3. *noise traders* make forecasts on the basis of data which they believe constitutes a signal, but is in fact uncorrelated with the future value of the asset [2][10].

Although chartist strategies should not be profitable according to the efficient markets hypothesis [8], this is not necessarily true if the market is outside of an efficient equilibrium. For example, if many agents adopt a chartist forecasting strategy it may be rational to follow suit as the chartist expectations may lead to a self-fulfilling prophecy in the form of a speculative bubble. Thus there are feedback effects from these three classes of forecasting strategy and it is important to study the interaction between them in order to understand the macroscopic behaviour of the market as a whole.

In the model the market mechanism is a continuous double auction with limit orders. This model is implemented as a discrete-event simulation using a Bernoulli process [1] to model time; on any given time-step, agents arrive at the market with probability λ.

Orders are executed using a time priority rule: the transaction price is the price of the order which was submitted first regardless of whether it is a bid or ask. If an order cannot be executed immediately it is queued on the order-book.

All orders have a limited order life, after which they are removed from the order book if they have not been successfully matched (a constant exogenously set to 200 units of time).

If a bid exceeds the best ask (lowest ask price on the order book), it is entered at the ask price (converted into a market order rather than a limit order). If an ask is lower than the best bid (highest bid price on the order book) it is entered at the bid price (again converted into a market order rather than a limit order).

The sign (buy or sell) and the price of the order for agent i at time t is determined as a function of each agent's *forecast* of the expected return $\hat{r}_{(i,t,t+\tau)}$ for the period $t + \tau$ (τ a constant defining the time horizon over which price expectations are made). The forecasted price for agent i is set according to:

$$p_{(i,t+\tau)} = p_t \cdot e^{\hat{r}_{(i,t,t+\tau)}} \tag{1}$$

where p_t is the quoted price at time t, and the sign of the order is *buy* iff. $p_{(i,t+\tau)} \geq p_t$ or *sell* iff. $p_{(i,t+\tau)} < p_t$.

The forecasted expected return for the period $t+\tau$ of agent i at time t is calculated with a linear combination of fundamentalist, chartist and noise-trader forecasting rules:

$$\hat{r}_{(i,t,t+\tau)} = \hat{r}_{f(i,t,t+\tau)} + \hat{r}_{c(i,t,t+\tau)} + \hat{r}_{n(i,t,t+\tau)} \tag{2}$$

$$\hat{r}_{f(i,t,t+\tau)} = f_{(i,t)} \cdot (\frac{F - p_t}{p_t}) \tag{3}$$

$$\hat{r}_{c(i,t,t+\tau)} = c_{(i,t)} \cdot \hat{r}_{L_i} \tag{4}$$

$$\hat{r}_{n(i,t,t+\tau)} = n_{(i,t)} \cdot \epsilon_{(i,t)} \tag{5}$$

In Equation 3, F is the "fundamental price" (which is exogenous and constant for all agents), p_t is given the value of the transaction at the previous time step or in the absence of a transaction the midpoint of the spread, $\epsilon_{(i,t)}$ are random i.i.d. variables distributed $\sim N(0,1)$ and r_{L_i} is a forecast based on historical data, in our case a moving average of actual returns over the horizon period L_i:

$$\hat{r}_{L_i} = \frac{1}{L_i} \sum_{j=1}^{L_i} \frac{p_{t-j} - p_{t-j-1}}{p_{t-j-1}} \tag{6}$$

The period L_i is randomly and uniformly initialised from the interval $(1, l_{max})$. The linear coefficients $f_{(i,t)}$, $c_{(i,t)}$ and $n_{(i,t)}$ denote the weight that agent i gives to each class of forecast amongst fundamentalist, chartist and noise-trader respectively at time t. Bids (b_t^i) and asks (a_t^i) (that is buys and sells) are entered into the market with a markup or markdown:

$$b_t^i = p_{t,t+\tau}^i (1 - k_i) \tag{7}$$

$$a_t^i = p_{t,t+\tau}^i (1 + k_i) \tag{8}$$

where k_i is randomly and uniformly initialised from an interval $(0, k_{max})$.

The *initial* values at time $t=0$ for the fundamentalist $f_{(i,0)}$, chartist $c_{(i,0)}$ and noise $n_{(i,0)}$ coefficients are drawn from the following distributions:

$$\begin{aligned} f_{(i,0)} &\sim |N(0,\sigma_f)|, \\ c_{(i,0)} &\sim N(0,\sigma_c), \\ n_{(i,0)} &\sim |N(0,\sigma_n)| \end{aligned} \tag{9}$$

Chiarella and Iori [2] study the impact of heterogeneous forecasting rules all of whom have varying degrees of chartist, fundamental and noise contributions. They found the existence of chartists generate large price jumps and realistic levels and dynamics of volatility of returns. In this model just a single stock is traded and all trades are simplified to have a volume of one. In the following, we refer to this model as the "CI Model". We study the characteristics of the model in the following sections, where we seek to understand the impact of different strategies on the distribution of returns.

3 Methodology and Treatments

We simulate the "CI Model" under various treatments described below. For each simulation, we randomly sample the values of free parameters, and record the moments of the resulting return distribution for later analysis. The default experiment time is 2.5×10^5 time units. Each experiment is executed for 3.5×10^5 units of time; the first 1.0×10^5 units of time are discarded to allow the model to stabilise. The experiment is repeated 3000 times with different randomly chosen parameter values.

3.1 Chiarella Iori Model Treatment

In our first experiment, we execute the "CI Model" without any modifications to try to understand what is the relationship between agent strategies and return distribution kurtosis in this model.

3.2 Trend Followers Only Treatment

One feature of the chartist trader is that it is both trend following and counter-trend following. In Equation 9, it can be seen that the chartist weight value can be either negative or positive. A positive weight means the agent expectation is a movement in the direction of the trend, while a negative weight means the agent expectation is counter to the direction of the trend (this is a contrarian behaviour). In this model, the chartist component represents two concepts which conflict with each other (trend following and counter trend following). We conjecture that the two behaviours may tend to cancel out phenomena like excess kurtosis. To test this conjecture, we modify the model such that we have only positive chartist weight (only trend following behaviour). Thus, Equation 9 becomes:

$$\begin{aligned} f_{(i,0)} &\sim |N(0, \sigma_f)|, \\ c_{(i,0)} &\sim |N(0, \sigma_c)|, \\ n_{(i,0)} &\sim |N(0, \sigma_n)| \end{aligned} \tag{10}$$

3.3 Sampled Moving Average Treatment

One feature of the CI Model is the weakness of the moving average (the chartist trend) calculation. The calculation introduces a large number of zero changes into the calculation of the trend which tends to weaken the trend, in addition any noise in the trend is also incorporated into the calculation. We modify the calculation made by agents by using sampling of transaction prices over varying sample sizes and varying horizons. This reflects the behaviour of actual traders, who employ such techniques to reduce noise in trend signals which allows them to better understand underlying trend changes.

3.4 Gaussian Noise Trader Treatment

The noise-trader forecast is obtained by multiplying the agent's strategy coefficient with a Gaussian random variate (Equation 5). However, the strategy coefficient is *itself* a Gaussian random-variate, and the product of two Gaussian variates is non-Gaussian. In addition, the employment of a markup/markdown distorts the distribution from which order prices are drawn, and hence, the markup/markdown is removed. We change Equation 5 to Equation 11 and pass Gaussian distributed order prices around the current market price to the market.

$$\hat{r}_{n(i,t,t+\tau)} = n_{(i,t)} \tag{11}$$

$$n_{(i,0)} \sim N(0, \sigma_n) \tag{12}$$

Gaussian Noise Traders Only Firstly we repeat the above simulation but only execute with noise traders (in Equation 2 $\hat{r}_{f(i,t,t+\tau)}$ and $\hat{r}_{c(i,t,t+\tau)}$ are set to zero).

All Traders Secondly, we repeat the experiment, but with traders adopting one of fundamentalist, chartist or noise strategies.

4 Discussion

We found with the Monte-Carlo execution of the original, unmodified Chiarella and Iori model [2] (under free-parameter variation) that chartist agents had no impact on the degree of return distribution kurtosis. The chartist in the original model embodied two different properties; a trend following, and a counter-trend following (a contrarian) property. We conjectured that the combination of the two properties suppressed the affects of trend following on the generation of non-Gaussian return distributions. We repeated the experiment with a model in which the chartist followed trends only, and we found that the trend following chartists did, indeed, have an impact on the kurtosis of the return distribution. With an increase in weighting of the trend following chartist contribution we recorded an increase in excess kurtosis.

We also noticed that, when we increased the size of the noise distribution from which the noise coefficients where drawn, we saw a resultant increase in kurtosis. A larger distribution would generate larger valued coefficients which, we suspected, would increase the randomness of the model, and therefore, make the resultant return distribution tend towards a Gaussian distribution. However, increasing the noise contribution increased the excess kurtosis of the return distribution; this seemed counter-intuitive. We re-configured the CI Model to run with just noise traders in a zero-intelligence treatment (in Equation 2 $\hat{r}_{f(i,t,t+\tau)}$ and $\hat{r}_{c(i,t,t+\tau)}$ are set to zero), the results produced a non-Gaussian return distribution. On examining the model, we noticed that the noise trader forecast is obtained by multiplying the agent's strategy coefficient with a Gaussian random variate (Equation 5). The noise trader coefficient is itself drawn from a Gaussian distribution, and the product of two Gaussian distributions is not a Gaussian distribution. In addition, the model was also performing a markup/markdown on the order prices of the noise trader. The combined effect of these model assumptions, ensured that the noise traders were entering orders with prices that were non-Gaussian distributed around the current market price. By modifying the noise trader order price expectation calculation such that the orders entered by the noise traders were Gaussian distributed about the current market price the model generated Gaussian return distributions. Importantly, this model, when configured to run with zero-intelligent traders, will generate Gaussian return distributions. The market micro structure has no impact on the return distribution, and it is, therefore, the "intelligent" traders behaviour which is responsible for the non-Gaussian return distributions. Other modellers have suggested that the market micro structure is responsible for non-Gaussian return distributions. In the Maslov model [9], for example, zero-intelligent agents introduce orders to the market with prices that are randomly drawn from a uniform distribution rather than a Gaussian distribution, it seems by construction that their model will generate non-Gaussian return distributions. The original unmodified Chiarella and Iori model [2] has also adopted model assumptions which by construction will generate non-Gaussian return distributions.

We have found that excess kurtosis is directly related to the relative weighting of fundamentalist, trend following and noise trading strategies. By adding the fundamentalist and chartist contributions to the zero-intelligence version of the CI model, we found that we no longer generated a Gaussian return distribution. With extensions to a simple model, we have shown that fundamentalist and chartist strategies impact return kurtosis within these models. Both the fundamentalist and chartist strategies look to past behaviour in the market; the fundamentalist generates orders based on the last behaviour of an agent in the market and the distance between the new price and a fundamental price, the chartist makes predictions based on price trends generated by other agents in the market. Both these contributions combine to increase the kurtosis of returns in the market. The chartist tending to increase the fat tails of the return distribution while the fundamentalist tending to increase the peak of the return distribution.

References

1. Bertsekas, D. and Tsitsiklis, J. (2008). *Introduction to Probability*. Athena Scientific.
2. Chiarella, C. and Iori, G. (2002). A Simulation Analysis of the Microstructure of Double Auction Markets. *Quantitative Finance*, 2:346–353.
3. Chiarella, C., Iori, G., and Perello, J. (2009). The impact of heterogeneous trading rules on the limit order book and order flows. *Journal of Dynamics and Control*, 33:525–537.
4. Gode, D. and Sunder, S. (1993). Allocative Efficiency of Markets with Zero-Intelligence Traders: Market as a Partial Substitute for Individual Rationality. *Journal of Political Economy*, 101(1):119–137.
5. Hommes, C. (2005). Coordination of Expectations in Asset Pricing Experiments. *Review of Financial Studies*, 18(3):955–980.
6. LeBaron, B. and Yamamoto, R. (2007). Long-Memory in an Order-Driven Market. *Physica A*, 383:85–89.
7. LeBaron, B. and Yamamoto, R. (2008). The Impact of Imitation on Long-Memory in an Order Driven Market. *Eastern Economic Journal*, 34:504–517.
8. Lo, A. (2005). Reconciling Efficient Markets with Behavioural Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting*, 7(2):21–44.
9. Maslov, S. (2000). Simple model of a limit order-driven market. *Physica A:Statistical Mechanics and its Applications*, 278:571–578.
10. Long, J. D., Shleifer, A., Summers, L., and Waldmann, R. (1990). Noise Trader Risk in Financial Markets. *The Journal of Political Economy*, 98(4):703–738.
11. Rayner, N., Phelps, S., and Constantinou, N. (2014). Learning is Neither Sufficient nor Necessary: An Agent Based Model of Long Memory in Financial Markets. *AI Communications*, 27: 437–452.
12. Yamamoto, R. and LeBaron, B. (2010). Order-splitting and Long-Memory in an Order-Driven Market. *European Physical Journal B*, 73(1):51–57.