AI Communication, Special Track Soft Computing in Finance and Economics
Accepted for publication, June 2013

Learning is Neither Sufficient Nor Necessary: An Agent-Based Model of Long Memory in Financial Markets

Neil Rayner [a,*], Steve Phelps [a]

[a] *Centre for Computational Finance and Economic Agents (CCFEA), University of Essex, Colchester, CO4 3SQ, UK.*
E-mail: {njwray,sphelps}@essex.ac.uk

Nick Constantinou [b]

[b] *Essex Business School (EBS), University of Essex, Colchester, CO4 3SQ, UK.*
E-mail: {nconst}@essex.ac.uk

Financial markets exhibit long memory phenomena; certain actions in the market have a persistent influence on market behaviour over time. It has been conjectured that this persistence is caused by social learning; traders imitate successful strategies and discard poorly performing ones. We test this conjecture with an existing adaptive agent-based model, and we note that the robustness of the model is directly related to the dynamics of learning. Models in which learning converges to a stationary steady state fail to produce realistic time series data. In contrast, models in which learning leads to continuous *dynamic* strategy switching behaviour in the steady state are able to reproduce the long memory phenomena over time. We demonstrate that a model which incorporates contrarian trading strategies results in more dynamic behaviour in steady state, and hence is able to produce more realistic results. We also demonstrate that a non-learning contrarian model that performs dynamic strategy switching produces long memory phenomena and therefore that learning is not necessary.

Keywords: long memory, agent-based models, stylised facts, contrarian, adaptive expectations

*Corresponding author: First Author, address of first author.

AI Communications
ISSN 0921-7126, IOS Press. All rights reserved

1. Introduction

In economics and finance agent-based modelling has become increasingly popular [10]. These agent-based models hold great promise in the construction of regulatory frameworks and financial market institutions that are robust within an uncertain world. They enable us to perform stress tests which incorporate effects such as behavioural biases or nonlinear emergent phenomena. Such models have been used, for example, to analyse the possible effects of different financial regulatory frameworks such as Basel II [12], Tobin taxes [27] and the effects of intervention by central banks on exchange rate fluctuations [38].

Agent-based models are able to capture complex trading behaviour and market micro-structure which is difficult to incorporate in traditional equation-based models. The rules that trading agents use to make their decisions can take the form of inductive heuristics which are gradually learnt over time [23]. This allows for more realistic models of trading agent decision making, using heuristics which are consistent with empirical observations from the controlled study of actual human subjects [13].

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. We focus on well known stylized facts of high-frequency time series data observed in real financial markets and we analyse to what extent different model assumptions are consistent with these phenomena.

In this research we test the hypothesis that the stylized facts of observed long memory are caused by agents imitating each other. Specifically, we analyse long memory properties with the following attributes:

1. over periods of time volume can be consistently high or low [25];
2. similar volatilities appear in the market in clusters [8,15,28,29,31];
3. signs of orders (that is, buy orders have a positive sign and sell orders have a negative sign) like volatility and volume exhibit long memory [3,19,20]; and
4. returns do not exhibit long memory [7], similar returns do not cluster together, and high frequency returns exhibit anti-persistence [32].

The outline of this paper is as follows. In Section 2 we provide some background on coevolutionary algorithms and their role in agent-based models. In Section 3 we describe an existing agent-based model [17,18] within which agents learn trading strategies from each other. We introduce our extensions to this important model, describing a model which possesses increased contrarianism. Contrarian traders trade against the perceived wisdom of the market and will for example sell when the market is buying. We analyse empirical data and validate the existing model in Section 4 and present our results in Section 5 where we show that this model does indeed generate long memory phenomena while learning. In Section 6 we analyse the ability of the two models to generate stable long memory phenomena under free-parameter variation and over extended periods of time. We demonstrate that the existing model is not robust to changes in free parameters and that in steady state the model does not produce, in general, long memory phenomena. However in contrast, our extended model does. We also perform a sensitivity analysis comparing agent strategy characteristics and their impact on long memory (following the approach of [39]). In Section 7 we run a non-learning version of our extended model reporting our results and demonstrating that long memory phenomena can be generated without learning. Finally in Section 8 we conclude.

2. Agent-Based models

Research in multi-agent systems has been influenced by theoretical models from economics and game-theory where agents behave with rational expectations [34]. Sometimes these theoretical models fail to capture observations in empirical financial data from real-world financial exchanges [24]. A. Lo [21] introduces the "adaptive markets hypothesis" where agents learn behaviours rather than adopting rational expectations. He suggests this may be able to explain observed phenomena present in empirical financial data.

Agent-based models simulate the microstructure and behaviour of the agents in a market with simple adaptive behaviours. B. Lebaron et al [17,16] adopt this approach and demonstrate that they are able to reproduce statistical qualities of empirical financial data ("stylised facts").

Agent-based computational finance (ACF) models can be usefully divided into two types [4,14]: firstly *N-type* models which adopt a limited set of agent types meaning the degree of heterogeneity can be low. For example a 2-type model might have two types of agent one employing a fundamentalist strategy and the other a chartist strategy (more will said about these strategy types in Section 3), and secondly *Santa-Fe Institute* models where each agent can be distinct from all other agents within the model. The agents within these models adapt strategies over time and tend to have a high degree of heterogeneity. Our research adopts the Santa-Fe Institute model approach and builds on work in [18].

M. Kampouridis et al [14] claim that financial markets are non-stationary, and that strategies that worked at one time become either ineffective or have reduced effectiveness over time. Adaption of strategies is important in maintaining strategy effectiveness over time. Researchers adopt various adaptive approaches for both the *N-type* and the *Santa-Fe Institute* types of model [14,18,30,33].

An important question in this research is how do agents interact and learn. Ecologists understand that when animals interact they can affect each others evolution e.g. predators and prey, or hosts and parasites. This interaction is called co-evolution where the fitness of individuals in the ecology is dependent on one another. Co-evolutionary algorithms can be interpreted as models of social learning [37] in which agents imitate strategies of other more successful agents. B. LeBaron et al [16] motivates the importance of coevolving models: "A trader's performance depends critically on the behaviour of others." M.

Kampouridis et al [14], S. Martinez-Jaramillo et al [30] and ourselves consider coevolution to be an important feature of financial markets, although we demonstrate in this paper that many of the "stylised facts" of financial markets can be produced without co-evolution. S. Martinez-Jaramillo et al [30] describe a coevolutionary development platform called CHASM which supports the population of a double auction order driven market for multiple heterogeneous agents. This research follows a similar approach to CHASM. However, S. Martinez-Jaramillo et al [30] use a genetic programming tool called EDDIE [35,36] to adapt strategies using a large variety of financial indicators while this research uses a much simpler model, adopting a genetic algorithm and a simpler set of financial indicators.

3. The Model

3.1. Agent Expectations

We adopt an adaptive expectations model with three classes of strategy which are used to form expectations about future returns:

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 [5,26].

Although chartist strategies should not be profitable according to the efficient markets hypothesis [21], 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.

We model the market mechanism as a continuous double auction with limit orders. Our model is implemented as a discrete-event simulation using a Bernoulli process [2] 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 [17,18].

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:

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

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

We adopt the framework of [17,18] in which 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 iid. 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:

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

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

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

3.2. *Evolutionary Algorithm*

To model imitation [17,18], agents use a co-evolutionary Genetic Algorithm to learn the coefficients $f_{(i,t)}$, $c_{(i,t)}$ and $n_{(i,t)}$. Each agent records its own forecast error as the market progresses and generates a fitness score s_i.

$$s_i = \frac{1}{\sum_{t=1}^{5000} (p_t - E_i(p_t))^2} \tag{9}$$

After every 5000 time steps, the population of agents reproduces. The fitness score of agent i is given by:

$$S_i = \frac{s_i}{\sum_i s_i} \tag{10}$$

We use these fitness scores to produce the next generation of agents using fitness proportionate selection [1]. The next generation consists of the same number of agents as the previous generation (steady-state evolution). Each member of the new generation is picked at random from the previous generation with a probability of S_i. When an agent is selected for reproduction, one of its strategy weights f, c, n or history horizon value L_i is inherited.

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

$$\begin{aligned} f_{(i,0)} &\sim |N(0, \sigma_f)|, \\ c_{(i,0)} &\sim N(0, \sigma_c), \\ n_{(i,0)} &\sim |N(0, \sigma_n)| \end{aligned} \tag{11}$$

In addition to the learning of weights after each 5000 units of time, agents mutate one of their weights with a probability p_m drawing a weight at random from the distributions in Equation 11 or drawing a new history horizon value (L_i).

3.3. *Contrarian Model*

We analyse two variants of this basic model; an existing model in the literature [17,18] in which forecasting strategies are linear *combinations* as per Equation 2 (henceforth we refer to this model as the LY Model), and our own model in which each agent adopts either an *atomic* fundamentalist (Equation 3), chartist (Equation 4) or noise trader (Equation 5) forecasting rule and not a linear combination as in the LY model above (Equation 2). Both our model and the LY model occupy the same strategy space. In our model however, two out of the three weights are zero reducing each agent to just one of the return forecast rules (Equation 3, Equation 4 or Equation 5).

J. Conrad et al [6] identify two main trading strategies: momentum strategies based on the following of trends, and contrarian strategies based on the reversal of trends. Contrarian traders predict price reversals and make profit (when they are correct) by positioning themselves to take advantage of that reversal. These two diametrically opposed strategies appear to exist simultaneously [6]. For example, in a rising trend momentum strategists will place bid orders (buys) in the market while contrarian strategists will place asks orders (sells). Other factors contribute to contrarian like behaviour. For example "pairs trading" [11] where a pair of related stocks are traded together such that when one is relatively expensive and the other cheap, traders sell the expensive one and buy the cheap one, a behaviour which is entirely independent of trends in the market. Other contrarian like behaviour can be caused by events external to the market, like arbitrage (risk-free profit) opportunities. For example, prices for the same stock can differ between two markets so traders buy in the cheaper market and sell in the more expensive market (gaining risk-free profit).

The LY Model implements contrarianism by allowing the chartist weight to go negative (Equation 11). In the LY Model agents imitate successful strategies, if the most successful strategy employs a negative chartist weight then agents will tend to adopt that strategy and this will represent the strategy of the herd. Contrarian strategists seek to behave in a contrary way to the herd. To capture this behaviour we introduce additional contrarianism. We add two contrarian strategies; one is to negate the learnt trend, so the contrarian agent predicts a price move in the opposite direction of the learnt trend (e.g if the learnt chartist trend is negative then the contrarians will predict a price move in a positive direction and vice versa), and the second that the price will not trend in the learnt direction at all but will remain at its current level. We have chosen to implement this firstly by setting the contrarian chartist strategy to be the negative of the non-contrarian chartist strategy:

$$\hat{r}_{c_c(i,t,t+\tau)} = -\hat{r}_{c(i,t,t+\tau)} \tag{12}$$

Secondly we set the contrarian fundamentalist and noise strategies to be the zeroed non-contrarian fundamentalist and noise strategies:

$$\hat{r}_{f_c(i,t,t+\tau)} = \hat{r}_{n_c(i,t,t+\tau)} = 0 \tag{13}$$

In the *contrarian* variant, agents can choose from the following discrete set of return forecasting strategies:

$$\{\hat{r}_{c(i,t,t+\tau)}, \hat{r}_{f(i,t,t+\tau)}, \hat{r}_{n(i,t,t+\tau)},$$
$$\hat{r}_{f_c(i,t,t+\tau)}, \hat{r}_{n_c(i,t,t+\tau)}, \hat{r}_{c_c(i,t,t+\tau)}\}$$

The same learning process operates in this model as in the LY model (but with the addition of the contrarian parameter). So an agent can change from fundamentalist to chartist or contrarian to non-contrarian to take advantage of a better strategy. During initialisation of the model values are drawn randomly from the distributions in Equation 11 as in the LY model, but each agent also chooses randomly between being a fundamentalist, chartist or noise trader and contrarian or non-contrarian. Henceforth we refer to this latter model as “the Contrarian Model”.

4. Methodology and Empirical Data

We compare model assumptions according to how well a particular model reproduces the long-memory “stylized facts” of volume, volatility, order signs, and returns. To compare models we test their long-memory properties using Lo’s modified rescaled range (R/S) statistic [22] (sometimes called range over standard deviation). The statistic is designed to compare the maximum and minimum values of running sums of deviations from the sample mean, re-normalised by the sample standard deviation. The deviations are greater in the presence of long-memory than in the absence of long-memory. The Lo R/S statistic includes the weighted auto-covariance up to lag q to capture the effects of short-range dependence. Firstly we examine the long-memory properties of real-world data.

The companies in the FTSE 100 represent approximately 80% of the capital value of the London Stock Exchange (LSE), and the FTSE 100 index is a leading UK market indicator. We examine data from major FTSE 100 companies; intraday data from April 2008 to June 2008 (the stocks presented in Table 1). The LSE operates an electronic double-auction market and historical intraday data can be purchased from the exchange. Data is provided in the form of three relational database tables from the LSE electronic marketplace SETS. Given these three database tables it is possible to infer the actual intraday behaviour of the market. We have analysed each stock in Table 1 for each day over the April 2008 to June 2008 period and present long memory results in Table 2. The Table shows the percentage of days of stock data that exhibit long memory for volume, volatility, order signs and returns. We see that we obtain long memory in volume, volatility, order signs but almost none in returns. Table 3 shows the average return kurtosis, and average Hurst exponents for volume, volatility, order signs and returns for all the stock days. Hurst values above 0.5 indicate persistence (the reinforcement of trends) while below 0.5 indicates anti-persistence (mean-revertion). All of volume, volatility, order signs are persistent while returns are anti-persistent. The results demonstrate the “stylized facts” present in real-world data confirming observations by other researchers as discussed in Section 1.

Table 1

Major FTSE 100 companies data set ranging over the period from April to June 2008. (The LSE dataset used is 3.42 GB in size and covers over 550 days)

Company	**Sector**
Astrazeneca Plc	Pharmaceuticals and Biotechnology
British American Tobacco Plc	Tobacco
BG Group	Oil and Gas Producers
Bhp Billiton Plc	Mining
British Petroleum Plc	Oil and Gas Producers
Bovis Homes Plc	Household Goods
Glaxosmithkline Plc	Pharmaceuticals and Biotechnology
HSBC Holdings PLC	Banks
Royal Dutch Shell Plc	Oil and Gas Producers

Table 2

Major FTSE 100 companies data. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×30 seconds to 10×30 seconds.

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	92.2	90.0	73.2	3.5
q=6	91.3	88.0	70.4	3.5
q=8	90.1	85.5	66.8	3.4
q=10	88.5	83.0	63.0	3.7

Table 3

Major FTSE 100 companies data (sampled every 30 seconds). Return distribution kurtosis. Hurst exponents for volume, volatility, order signs and returns.

Return Kurtosis	**Hurst Volume**	**Hurst Volatility**	**Hurst Order Signs**	**Hurst Returns**
16.9	0.70	0.66	0.66	0.46

Our first experiment tests the conjecture that imitation (in the LY Model) is sufficient to produce the long memory phenomena by attempting to reject the null hypothesis that long memory is caused by the strategies that each agent adopts, and not by learning at all. The model is simulated in two sequential phases: firstly, a learning phase, and secondly, a commitment phase. In the learning phase the agent's genetic algorithm searches for strategies with high relative fitness (see Equation 9 and Equation 10). In the commitment phase agent's commit to a learnt strategy and perform no further learning.

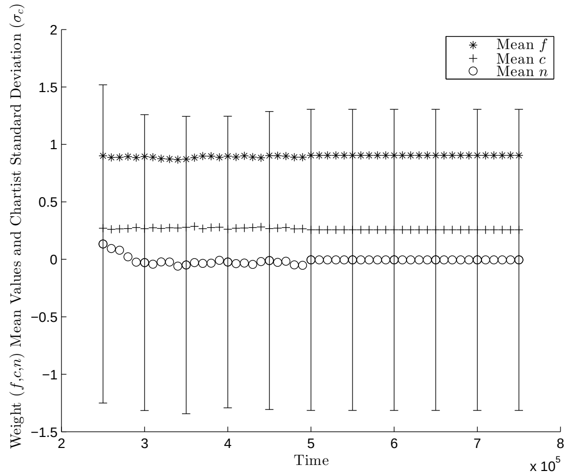

Fig. 1. LY Model. Fundamentalist (f), Chartist (c), Noise (n) mean weights and Chartist Standard Deviation (σ_c) (represented by the bars) with time. In the first half of the experiment learning is switched on and we see movements in mean weight values and changes in the standard deviations of weight distributions (only chartist shown here). In the second part of the experiment learning is switched off and the means of weights and weight distribution cease to change.

The default experiment time is 2.5×10^5 time units (consistent with [17,18]). The experiment is executed for three times the default experiment time (3 x 2.5×10^5 units of time); the first 2.5×10^5 units of time are discarded to allow the model to stabilise. The learning phase is executed for the remaining half of the experiment time (the default time) and then the commitment phase runs for the same period (default parameter values are displayed in Table 4 these values have been taken from [17,18]). Parameter values are varied randomly as described in Table 7 and Table 10.

5. Validation Results for LY Model

In Figure 1 we show the mean value (across all agents) of the fundamentalist (f), chartist (c) and the noise trader weight (n) with respect to time. It also shows the chartist weight distribution standard deviation (σ_c) (the noise and fundamentalist weight distribution standard deviation are not shown but behave in a similar manner). As we can see in Figure 1 there is a period of fluctuation in the mean weights as the agents move about in the strategy space. However, as we would expect, when the commitment phase starts at time $t = 5 \times 10^5$, the mean and standard deviation of the weights remains constant.

The experiment was executed 500 times. Tables 5 and 6 show the percentage of simulation runs which exhibit long memory in volume, volatility, signs of orders (buy or sell orders) and returns. In the first phase (the learning phase presented in Table 5) we see the long memory characteristics we are expecting with this model. In the second phase (the commitment phase presented in Table 6) long memory properties have largely disappeared. It is not, therefore, sufficient to have the correct mix of strategies in order to generate long memory. This implies that there is something about the dynamics of weight changing (caused in this case by the learning process) which is causing these phenomena.

One feature of the results obtained is the relative weakness of the long memory for order signs. These models make one particular simplifying assumption that may have implications for order sign long memory; the absence of an inventory, all trader orders are presented to the market with one unit of asset. In reality financial traders have a finite asset resource and present orders of varying volume sizes. In particular when traders need to buy or sell large quantities of an asset they will often split their orders into a set of smaller orders to minimise the market impact (the affect on asset price). A large single order entered into the market can attract very poor prices i.e a large order to buy will consume sellers orders starting with the cheapest and depending on the order size and the available volume on the order book will consume more and more expensive orders. F. Lillo et al [20] demonstrate that order-splitting introduces long-memory in order signs also B. LeBaron et al [16] demonstrate that the LY Model with order-splitting (and without imitation) produces order sign long memory. It is our conjecture that with the addition of an inventory and order splitting the weak order long memory results might improve.

Table 4
Default Values for All Models.

Parameter	**Value**
Std dev of fundamental weight (σ_f)	1.0
Std dev of chartist weight (σ_c)	1.5
Std dev of noise weight (σ_n)	0.5
Probability of mutation per generation (p_m)	0.08
Probability of entering the market λ per time step	0.5
Maximum markup or markdown (k_{max})	0.5
Maximum period over which trends are calculated (l_{max})	100 units of time
Period over which price expectations are made (τ)	200 units of time
Fundamental price (F)	1000
Number of Traders	1000
Order Life	200 units of time
Tick Size (the smallest price differential)	0.1

Table 5
LY Model Learning Phase. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	70	74	13	2
q=6	69	71	11	2
q=8	68	69	11	3
q=10	68	68	10	3

Table 6
LY Model Commitment. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	2	3	2	4
q=6	2	2	2	4
q=8	2	2	3	4
q=10	3	1	3	4

6. Model Robustness

In this section we review the robustness of the LY Model and the Contrarian Model by performing two sets of experiments. The two models differ from each other. The LY model adopts a lin-

Table 7

Ranges of Parameter Values

Parameter	**Value**
Probability of mutation per generation (p_m)	0.05 to 0.20
Probability of entering the market (λ)	0.5 to 1.0
Maximum markup or mark-down (k_{max})	0 to 0.5
Maximum period over which trends are calculated (l_{max})	50 to 150 units of time
Period over which price expectations are made (τ)	1 to 200 units of time

ear combination of rules whereas the Contrarian model uses an atomic forecast strategy specified by Equations 3-5. In addition, in the Contrarian model, agents adopt a contrarian or non-contrarian strategy. In the first experiment we do not vary free-parameters associated with the forecasting strategy and in the second we do. Any impact on robustness due to this difference is then apparent.

6.1. Free Parameter Variation

In this experiment we vary a selection of the free-parameters by drawing from uniform distributions described in Table 7. We have extended the experiment execution time to highlight any problems in robustness with respect to time. The experiments were run for 10 times the default time (10 x 2.5×10^5 units of time). Fifty sets of random parameter variations were executed with 10 executions for each set (totalling 500 for 2.5×10^6 units of time). The first 2.5×10^5 units of each experiment are discarded to allow the models time to stabilise.

In Table 8 we present the results of the first experiment. We note the LY Model produces weak order sign long memory. With the Contrarian Model (Table 9) the order sign long memory is about double that of the LY Model. In addition the percentage of runs with long memory in volatility and volume is significantly larger. The Contrarian Model is producing results consistent with the empirical data observed in Table 2 apart from a weakness in order sign long memory.

Table 8

Free Parameter Variation Experimental Results for LY Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	59	60	18	1
q=6	57	55	16	1
q=8	55	53	14	2
q=10	54	51	12	2

Table 9

Free Parameter Variation Experimental Results for Contrarian Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	85	78	37	0
q=6	84	77	37	0
q=8	83	77	38	0
q=10	82	76	38	0

6.2. Extended Free Parameter Variation

In the LY Model we sum contributions of chartist, fundamental and noise returns to make our return prediction (as in Equation 2) while with the Contrarian Model only one of chartist, fundamental and noise returns contributes to our prediction. In the first set of experiments we varied parameters that do not contribute to weights assigned to chartist, fundamental and noise returns demonstrating that we obtain improvements with the Contrarian Model without this difference. In the second set of experiments we vary some additional free-parameters; the standard deviations of the Gaussian distributions from which the weights are chosen (Equation 11). The experiments, as before, are run for 10 times the default time (10 x 2.5×10^5 units of time) and parameter values were randomly drawn (uniformly) from the ranges in Table 10 and 7. Fifty sets of random parameter variations were executed with 10 executions for each set (totalling 500 for 2.5×10^6 units of time). The first 2.5×10^5 units of each experiment are discarded to allow the models time to stabilise. The Contrarian Model is executed in two modes: firstly

Table 10

Ranges of Additional Parameter Values. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Parameter	**Value**
Std dev of fundamental weight (σ_f)	0.0 to 3.0
Std dev of chartist weight (σ_c)	0.0 to 3.0
Std dev of noise weight (σ_n)	0.0 to 3.0

a reduced mode where the contrarian functionality is switched off leaving just the atomic extension to the LY Model (this we refer to as the Atomic mode), and secondly the normal mode where the contrarian functionality is switched on and which includes the atomic extension to the LY Model. We do this to illustrate the contributions made by each extension.

In Table 11 we present the results from the LY Model. We note that the LY Model produces weak order sign long memory. While Table 12 and Table 13 show the improvements of the extensions introduced by the Atomic Mode and Contrarian Model. The Contrarian Model (Table 13) has strong long memory in volume and volatility, much improved order sign long memory. In Tables 14 and 16 we show the Hurst exponents for volume, volatility, order signs and returns, here we can see that we obtain a low Hurst exponent for returns of between 0.3 and 0.4 indicating anti-persistence which is consistent with observations of real world high frequency return data [32] and Table 3. We note also, that we get persistence in volume, volatility, and order signs (although persistence in order signs is weak under these variations of parameter settings). In these tables we also show the return distribution kurtosis, it is a stylized fact of real world markets that returns exhibit fat tails (a high kurtosis) [9]. Excess kurtosis may indeed be closely linked to volatility clustering and thus may be linked to long memory as well. We obtain for the two models values of 4.5 and 4.7 (a Gaussian distribution has a kurtosis of 3). These models, therefore, do produce fat tails (although not as strongly as the empirical observations in Table 3).

In Tables 17 and 18 we look at the first 2.5×10^5 units of time of an experiment. We have separated out the long memory properties of the execution of the LY Model into an early part of the test (the first 1.25×10^5 units of time) and a later part (the second 1.25×10^5 units of time). The experiment

Table 11

Extended Free Parameter Variation Experimental Results for LY Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	63	63	20	1
q=6	61	58	16	1
q=8	60	57	14	1
q=10	60	55	13	1

Table 12

Extended Free Parameter Variation Experimental. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time). Results for Atomic Mode

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	95	87	13	0
q=6	95	87	13	0
q=8	96	87	12	1
q=10	96	87	12	1

Table 13

Extended Free Parameter Variation Experimental Results for Contrarian Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	90	84	41	0
q=6	89	82	41	0
q=8	88	81	42	0
q=10	87	80	42	0

Table 14

Extended Free Parameter Variation Experimental Results for LY Model. Return distribution kurtosis. Hurst exponents for volume, volatility, order signs and returns.

Return Kurtosis	**Hurst Volume**	**Hurst Volatility**	**Hurst Order Signs**	**Hurst Returns**
4.7	0.59	0.61	0.54	0.4

Table 15

Extended Free Parameter Variation Experimental Results for Atomic Mode. Return distribution kurtosis. Hurst exponents for volume, volatility, order signs and returns.

Return Kurtosis	**Hurst Volume**	**Hurst Volatility**	**Hurst Order Signs**	**Hurst Returns**
4.3	0.64	0.66	0.50	0.30

Table 16

Extended Free Parameter Variation Experimental Results for Contrarian Model. Return distribution kurtosis. Hurst exponents for volume, volatility, order signs and returns.

Return Kurtosis	**Hurst Volume**	**Hurst Volatility**	**Hurst Order Signs**	**Hurst Returns**
4.5	0.66	0.64	0.52	0.31

Table 17

Early Phase Execution Results for LY Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time). The experiment time is for just the first 1.25×10^5 units of time of the default 2.5×10^5 units of time and not the full default 2.5×10^5 units of time.

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	100	99	50	50
q=6	100	99	52	55
q=8	100	99	53	57
q=10	100	99	54	61

Table 18

Later Phase Execution Results for LY Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time). The experiment time is for just the second 1.25×10^5 units of time of the default 2.5×10^5 units of time and not the full default 2.5×10^5 units of time.

Lag	**Volume**	**Volatility**	**Order Signs**	**Returns**
q=4	69	79	27	25
q=6	70	77	28	28
q=8	70	76	29	29
q=10	70	75	29	31

has been run with the default parameters in Table 4. We note that the long memory properties of the model are changing with respect to time. The

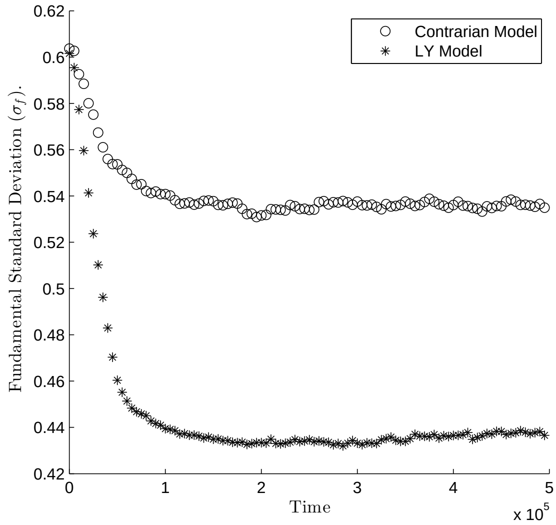

Fig. 2. LY and Contrarian Model fundamentalist standard deviation with time. The LY Model converges to a narrower distribution, while the Contrarian Model converges but to a much lesser extent.

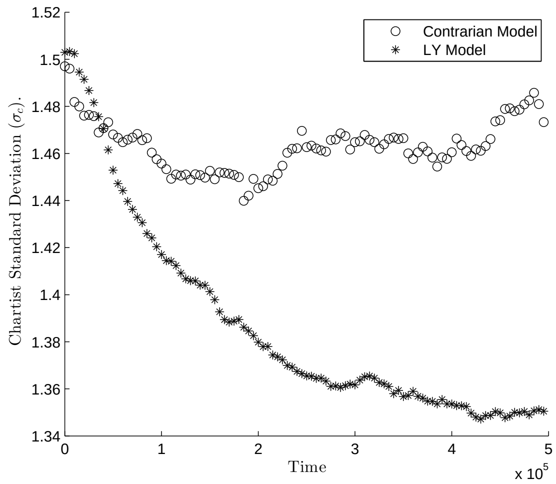

Fig. 3. LY and Contrarian Model chartist standard deviation with time. The LY Model converges to a narrower distribution, while the Contrarian Model although it fluctuates does not move far from it's initial state.

dynamics of the LY Model is changing the long memory properties of the model.

In order to appreciate why we have differences in the results of the two models we investigate the dynamic behaviour of the LY Model and contrast it with the Contrarian Model. During learn-

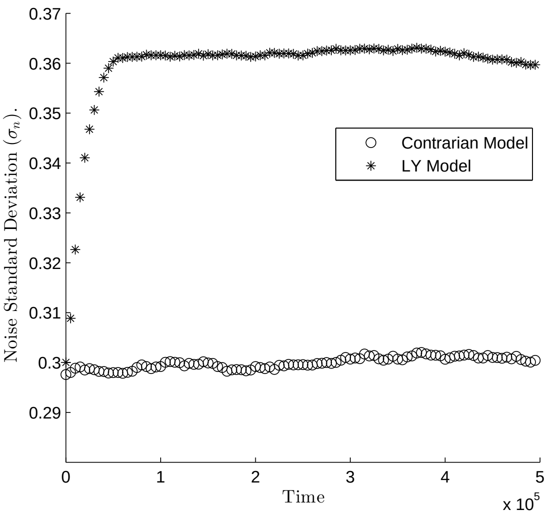

Fig. 4. LY and Contrarian Model Noise standard deviation with time. The LY Model diverges to a broader distribution, while the Contrarian Model remains close to constant.

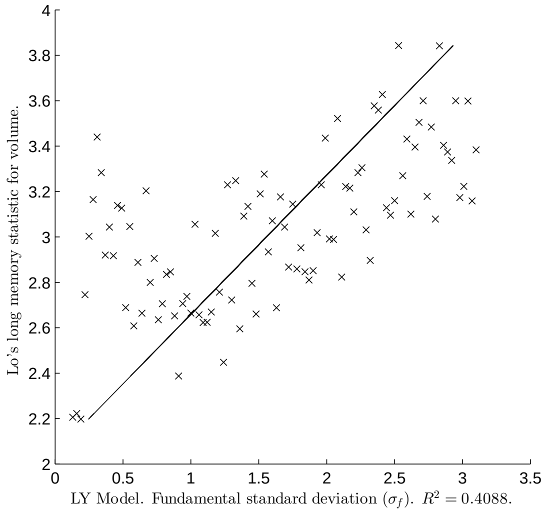

Fig. 5. LY Model. Long Memory in Volume Versus Fundamentalist Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the initial standard deviation of the fundamental weight distribution we note that we get decreasing long memory values in volume.

ing, the strategy weights will change as agents imitate more successful strategies. We run the LY Model and Contrarian Model 100 times using the default parameter set in Table 4. We present the results in Figures 2, 3 and 4. The figures show the standard deviations of the three strategy weight distributions (fundamental (σ_f), chartist (σ_c) and noise (σ_n)) with respect to time. In Figure 2 the changes in fundamendalist standard deviation (σ_f) for the LY Model and Contrarian Model are similar in that both converge from their initial state to a narrower distribution. However, the LY Model converges the most. In Figure 3 the changes in chartist standard deviation (σ_c) for the LY Model and Contrarian Model are quite different. The LY Model is clearly converging to a narrower distribution with the result that the agents are becoming increasingly more homogeneous. The Contrarian Model, although it is clearly changing, it's chartist weight distribution is not significantly converging away from it's initial state. In Figure 4 the changes in noise standard deviation (σ_n) for the LY Model and Contrarian Model are quite different again. The LY Model is clearly diverging to a broader distribution while the Contrarian Model is barely changing at all. The two Models are noticeably different in that the LY Model's weight distribution changes are much greater than those for the Contrarian Model. The dynamic characteristics of the Contrarian Model are more stable than the LY Model.

Given that the LY Model, in particular, has dynamics which are significantly changing the strategy weight distributions, we test to see what the effects on long memory are with changing the standard deviations of the weight distributions. We execute the models with the default parameter settings in Table 4 and with 100 different initial values of standard deviation in the range 0.13 to 3.1 and plot the results. We report the R^2 as a guide although many of the relationships are non-linear.

Figures 5 and 6 illustrate that as the fundamental standard deviation (σ_f) becomes narrower the long memory statistic in volume and volatility reduces. The convergence of the LY Model in Figure 2 is going to have a detrimental affect on the generation of long memory for volume and volatility. (We present a more detailed analysis of the effects on long memory through changes to strategy weight distribution in the Appendix)

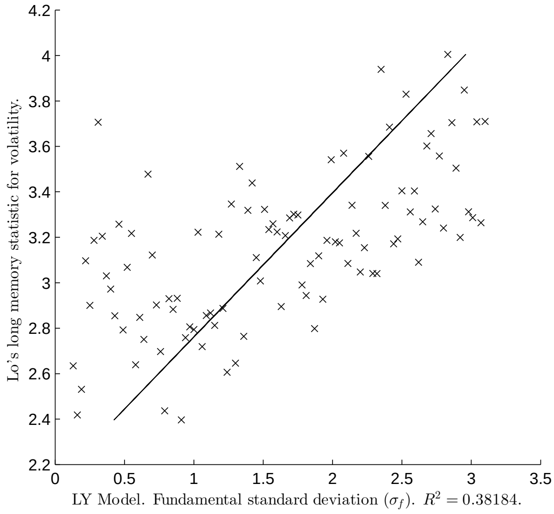

Fig. 6. LY Model. Long Memory in Volatility Versus Fundamental Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the initial standard deviation of the fundamental weight distribution we note that we get decreasing long memory values in volatility.

7. Non-Learning Model

In this section we test a simplified Contrarian model to illustrate the central role of the dynamics in producing long memory phenomena.

In this model we modify the Contrarian Model by switching off the learning process, but retaining mutation. Agents within this model do not imitate each other, there is no social learning, agents modify their strategy weights independently of each other through random mutation. The selection and cross-over are switched off, we still continue to introduce new strategies into the population via the mutation operator. Each agent is initialised with weights drawn from the distributions in Equation 11. After each 5000 units of time, agents may mutate one of their weights drawing a weight at random from the distributions in Equation 11 or drawing a new history horizon value (L_i) or switching from contrarian to non-contrarian (with a probability p_m). The model therefore has no converging learning process only the random mutation of the strategy space. The experiment is a repeat of the extended free-parameter variation experiment performed previously, but using a non-learning Contrarian Model.

By comparing Table 19 and Table 13 we can see there is very little difference in results between the learning Contrarian Model and the non-learning Contrarian Model. We conclude from this that learning or imitation is not necessary for the generation of long memory and that other models that have the dynamics of agent strategy change will also produce long memory. In Table 20 we show the Hurst exponents for volume, volatility, order signs and returns. We obtain a low Hurst exponent for returns of about 0.3 indicating anti-persistence again, as in Tables 14 and 16. In Table 20 we also show the return distribution kurtosis which demonstrates that this model also produces fat tails.

Table 19

Experimental Results for Non-Learning Contrarian Model. Percentages of runs with long memory for volume, volatility, order signs and returns at various time lags ranging from 4×50 (200 units of time) to 10×50 (500 units of time).

Lag	Volume	Volatility	Order Signs	Returns
q=4	89	83	40	0
q=6	88	82	41	0
q=8	87	81	41	0
q=10	86	80	42	0

Table 20

Extended Free Parameter Variation Experimental Results for Non-Learning Contrarian Model. Return distribution kurtosis. Hurst exponents for volume, volatility, order signs and returns.

Return Kurtosis	Hurst Volume	Hurst Volatility	Hurst Order Signs	Hurst Returns
4.5	0.67	0.64	0.52	0.31

8. Conclusion

While imitation may contribute to the generation of long memory phenomena in real financial markets other factors must play a role in the production of robust long memory phenomena over time (e.g contrarianism). M. Kampouridis et al [14] point towards a need for continuous co-evolution and adaption of strategies as a result of the non-stationary nature of real world financial markets. Over-commitment to small sets of strategies is ineffective over time. Our model which incorporates contrarianism and strong disparity between strategies is able to generate a more dynamic behaviour in steady state. In contrast the LY model is not as robust with respect to variation in free-parameter settings and execution time. This is caused by the convergence of the Genetic Algorithm to a smaller and smaller strategy space and a loss, therefore, of the dynamic that causes the long memory phenomena. We extended the important LY model adding atomic agents and increased contrarianism (the Contrarian Model) which retains the dynamic necessary to generate robust long memory phenomena (Section 6).

We demonstrated that it is indeed the dynamics of strategy changing that is the cause of long memory in these models. A simple Contrarian Model with no social learning and no imitation but with agents randomly and independently changing strategy generates long memory. Learning is not necessary for the generation of long memory.

Future Research

We conjecture that any model that causes and maintains a dynamic strategy switching behaviour will produce stable long memory and that non-learning heuristic models would also produce positive long memory. We intend to calibrate the discussed models to London Stock Exchange high frequency data with a view to identifying weaknesses in their ability to produce the empirical properties of the data.

The weakness we experience in long memory for order signs we conjecture is as a result of a simplified inventory in these models. [20] and [16] find that the introduction of order-splitting results in long memory in order signs. With the addition of order-splitting we may improve the order sign long memory characteristics of these models.

We expect to do further analysis of the two models including analysing a non-learning version of the LY Model.

Acknowledgements

The authors would like to thank the editor and the three anonymous referees of this paper, and the referees of the Robustness and Reliability of Electronic Marketplaces workshop, International Conference on Electronic Commerce (2011), all of whose comments helped clarify the work.

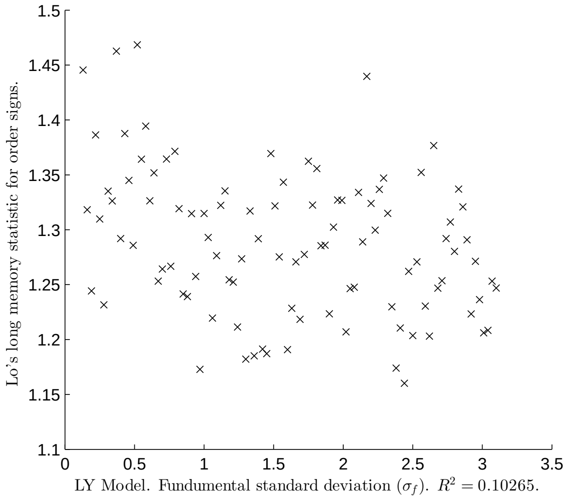

Fig. 7. LY Model. Long Memory in Order Sign Versus Fundamental Standard Deviation (σ_f) with lag ($q = 4$). As we can see there is little if any linear or non-linear relationship between order signs and size of the fundamental weight distribution.

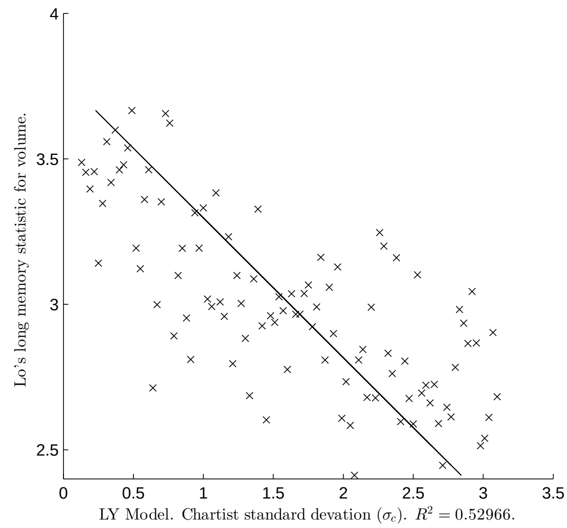

Fig. 8. LY Model. Long Memory in Volume Versus Chartist Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the increase the standard deviation of the chartist weight distribution we note that we get decreasing long memory values in volume.

Appendix

In this appendix we present a more detailed analysis of the effects on long memory through

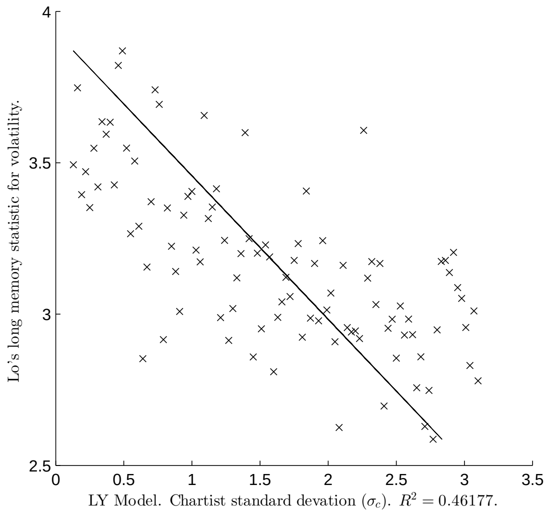

Fig. 9. LY Model. Long Memory in Volatility Versus Chartist Standard Deviation (σ_f) with lag ($q = 4$). As we increase the initial standard deviation of the fundamental weight distribution we note that we get decreasing long memory values in volatility.

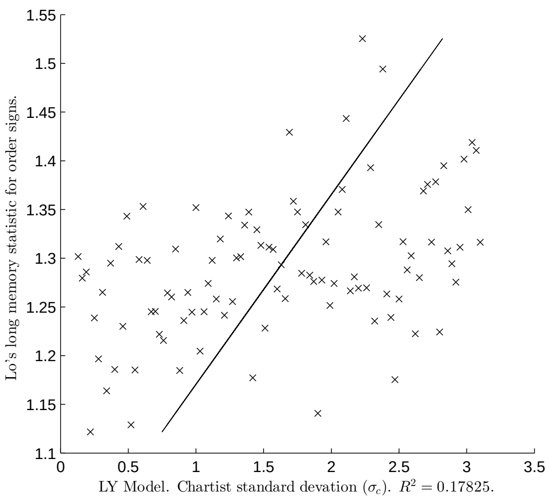

Fig. 10. LY Model. Long Memory in Order Signs Versus Chartist Standard Deviation (σ_f) with lag ($q = 4$). As we increase the initial standard deviation of the chartist weight distribution we note that we get decreasing long memory values in volume.

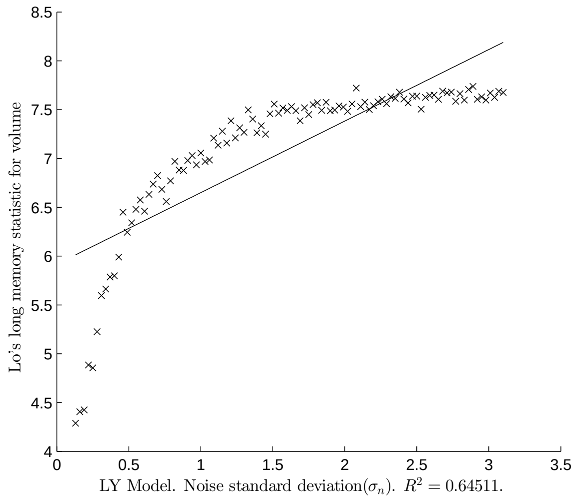

Fig. 11. LY Model. Long Memory in Volume Versus Noise Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the initial standard deviation of the noise weight distribution we note that we get decreasing long memory values in volume.

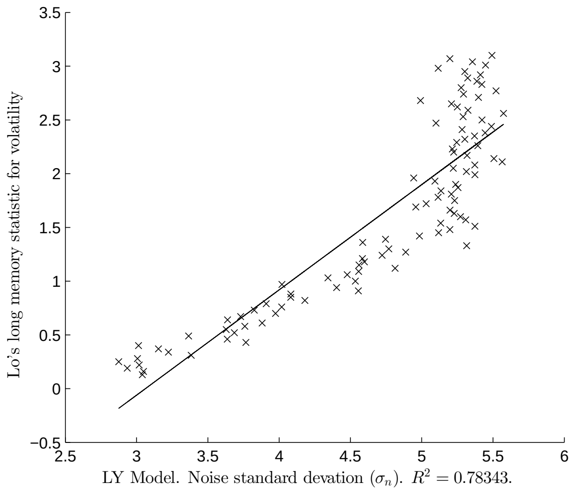

Fig. 12. LY Model. Long Memory in Volatility Versus Noise Standard Deviation (σ_f) with lag ($q = 4$). As we increase the initial standard deviation of the noise weight distribution we note that we get increasing long memory values in volatility.

changes to strategy weight distributions. In Figure 7 the relationship between the fundamentalist standard deviation (σ_f) and order sign long memory statistic is not significant for these runs. In Figures 8 and 9 we see the opposite relationship to Figures 5 and 6 with a convergence to a lower chartist standard deviation (σ_c) having a positive effect on the generation of long memory. In Figure 10 we have no significant effect on order sign long memory statistic. In Figures 11, 12 and 13

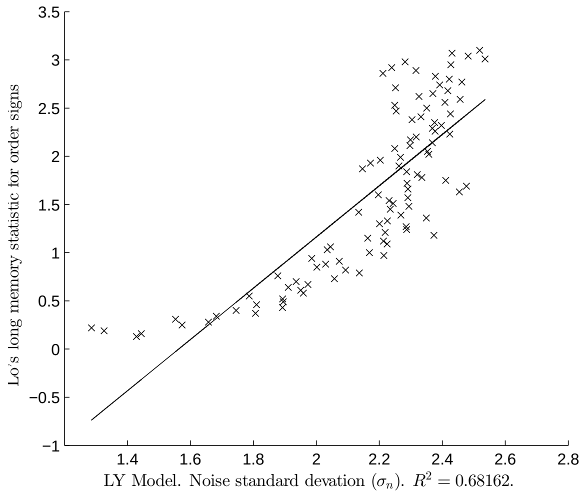

Fig. 13. LY Model. Long Memory in Order Signs Versus Fundamentalist Standard Deviation (σ_f) with lag ($q = 4$). As we increase the initial standard deviation of the noise weight distribution we note that we get increasing long memory values in order signs.

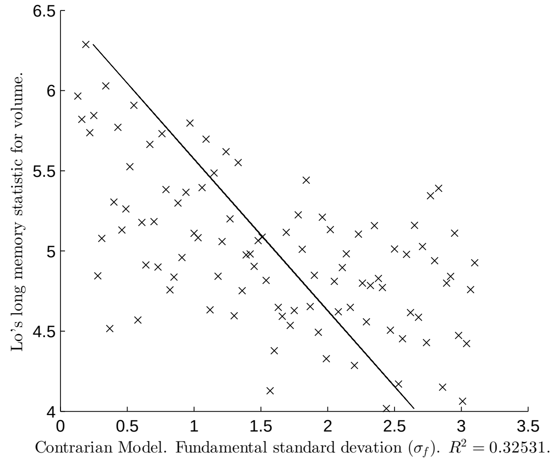

Fig. 14. Contrarian Model. Long Memory in Volume Versus Fundamentalist Standard Deviation (σ_f) with lag ($q = 4$). As we increase the initial standard deviation of the fundamental weight distribution we note that we get decreasing long memory values in volume.

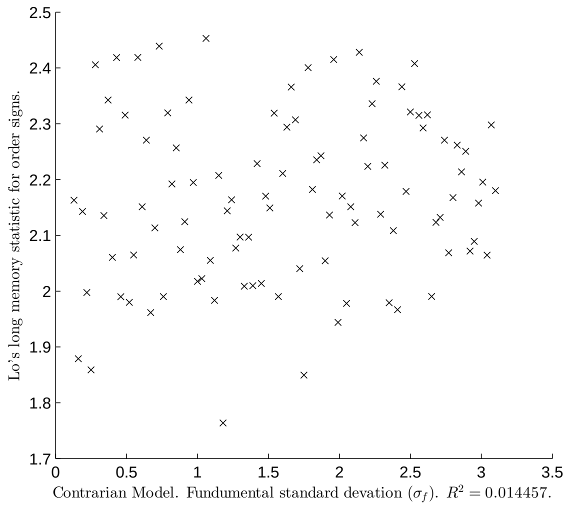

Fig. 15. Contrarian Model. Long Memory in Order Signs Versus Fundamentalist Standard Deviation (σ_f) with lag ($q = 4$).

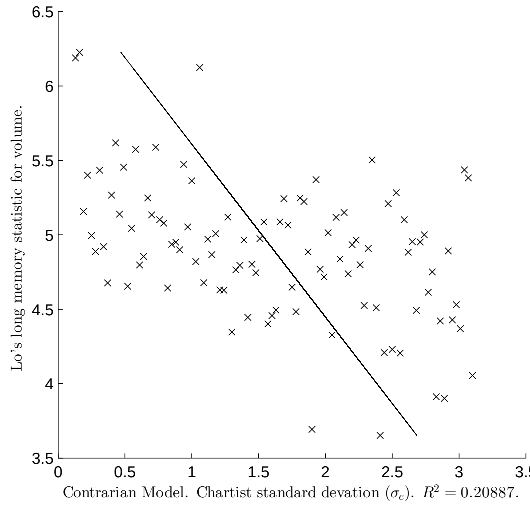

Fig. 16. Contrarian Model. Long Memory in Volume Versus Chartist Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the initial standard deviation of the chartist weight distribution we note that we get decreasing long memory values in volume.

volume, volatility and order sign long memory increase with an increase in noise standard deviation (σ_n).

To understand the reasons for the relationships between weight distributions and long memory it is important to understand that orders that are distant from the spread (the spread being the price difference between the best ask (best sell) and best bid (best buy) orders) are unlikely to transact or affect the mid price (the price halfway between the best bid and best ask). Volatility and volume are affected only by what happens in the region

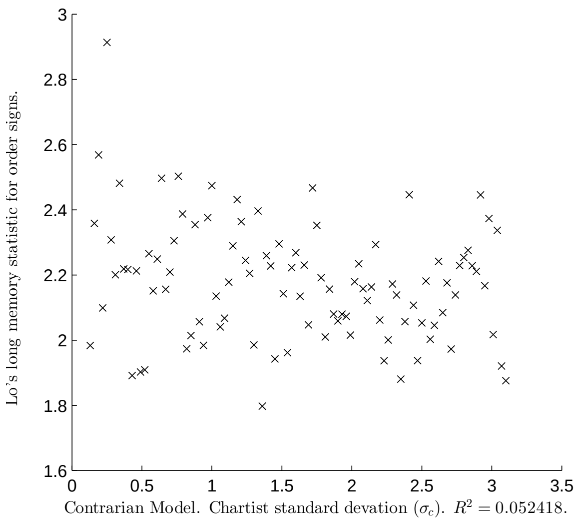

Fig. 17. Contrarian Model. Long Memory in Order Signs Versus Chartist Standard Deviation (σ_f) with lag ($q = 4$). As we decrease the initial standard deviation of the fundamental weight distribution we note that we get decreasing long memory values in volume.

of the spread. Large weight standard deviations will generate large weights so that orders are likely to be entered at large distances from the spread and therefore have no impact on the price in the market (transaction prices or mid prices). Small weight standard distributions will have the opposite effect.

The fundamentalist trader predicts a move back to a fundamental price and places orders that reflect that. The fundamentalist is mean-reverting and is always encouraging the market to return to the fundamental price. Mean-reversion is an anti-persistent behaviour and represses long memory. A small weight standard deviation for the fundamentalist means more anti-persistent behaviour around the spread and so a reduction in long memory. In Figures 5 and 6 we see this affect.

The noise trader is contributing a noise component to the market; drawing return predictions from a standard normal distribution which it then weights. If the noise standard deviation (σ_n) is small it will affect the spread introducing Gaussian random orders into the market which will affect the market price and repress long memory. In Figures 11, 12 and 13 we see that for small noise standard deviations long memory in volume, volatility and order signs is reduced.

The chartist trader follows trends, and so will generate orders which reflect past behaviour in the market. For small chartist standard deviations (σ_c) orders will be placed close to the spread. In Figures 8 and 9, with small values for the chartist standard deviation (σ_c), we observe stronger long memory. Given this sensitivity analysis, we suggest that it is the chartist trader which is primarily responsible for causing the long memory we detect in the LY Model.

In Figure 2 the fundamental standard deviation (σ_f) for the Contrarian Model does converge, though not as strongly as the LY Model. In Figure 14 we see that with a decrease in fundamental standard deviation we get larger long memory statistics for volume (the behaviour for volatility is similar). The convergence of the fundamental standard deviation (σ_f), in Figure 2, rather than decreasing long memory in volume and volatility is increasing it. Figure 15 shows that there is no relationship between the fundamental standard deviation and order sign long memory. In Figure 3 the chartist standard deviation (σ_f) for the Contrarian Model does not converge significantly consistent with Figure 16 where we see little relationship between the chartist standard deviation and long memory for volume (the behaviour for volatility is similar). Again, in Figure 17 there is no relationship between the chartist standard deviation and long memory for order sign.

The fundamentalist trader in the Contrarian Model has a contrarian component. While non-contrarian fundamentalists are still mean-reverting and therefore anti-persistent. Contrarian fundamentalists are not mean-reverting and are using the last market price.

References

[1] T. Back. *Evolutionary Algorithms in Theory and Practice.* Oxford Univ. Press, 1996.

[2] D.P. Bertsekas and J.N. Tsitsiklis. *Introduction to Probability.* Athena Scientific, 2008.

[3] J. Bouchaud, Y. Gefen, M. Potters, and M. Wyart. Fluctuations and Response in Financial Markets: The Subtle Nature of "random" price changes. *Quantitative Finance*, 4(2):176–190, 2004.

[4] S.-H. Chen, C.-L. Chang, and Y.-R. Du. Agent-Based Economic Models and Econometrics. *The Knowledge Engineering Review.*, 27(1):1–46, 2012.

[5] C. Chiarella and G. Iori. A Simulation Analysis of the Microstructure of Double Auction Markets. *Quantitative Finance.*, 2:346–353, 2002.

[6] J. Conrad and G. Kaul. An Anatomy of Trading Strategies. *The Review of Financial Studies.*, 11(3):489–519, 1998.

[7] R. Cont. Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues. *Quantitative Finance.*, 1:223–236, 2001.

[8] Z. Ding, C.W.J. Granger, and R.F. Engle. A Long Memory Property of Stock Market Returns and a New Model. *Journal of Empirical Finance.*, 1:83–106, 1993.

[9] R.F. Engle and A.J. Patton. What Good Is a Volatility Model? *Quantitative Finance.*, 1:237–245, 2001.

[10] J.D. Farmer and D. Foley. The Economy Needs Agent-Based Modelling. *Nature*, 460:685–686, 2009.

[11] E. Gatev, W.N. Goetzmann, and K.G. Rouwenhorst. Pairs Trading: Performance of a Relative-Value Arbitrage Rule. *Review of Financial Studies.*, 19(3):797–827, 2006.

[12] O. Hermsen. Does Basel II Destabilize Financial Markets? An Agent-based Financial Market Perspective. *The European Physical Journal B.*, 73(1):29–40, 2010.

[13] C. Hommes. Coordination of Expectations in Asset Pricing Experiments. *Review of Financial Studies.*, 18(3):955–980, 2005.

[14] M. Kampouridis, S.-H. Chen, and E. Tsang. Investigating the Effect of Different GP Algorithms on the Non-Stationary Behavior of Financial Markets. In *Proceedings of IEEE Symposium on Computational Intelligence for Financial Engineering and Economics.*, pages 43–50, 2011.

[15] C.G. Lamoureux and W.D. Lastrapes. Heteroskedasticity in Stock Return Data:Volume Versus GARCH Effects. *Journal of Finance.*, 45(1):221–229, 1990.

[16] B. LeBaron and R. Yamamoto. A Builders Guide to Agent Based Financial Markets. *Quantitative Finance*, 1:254–261, 2001.

[17] B. LeBaron and R. Yamamoto. Long-Memory in an Order-Driven Market. *Physica A.*, 383:85–89, 2007.

[18] B. LeBaron and R. Yamamoto. The Impact of Imitation on Long-Memory in an Order Driven Market. *Eastern Economic Journal.*, 34:504–517, 2008.

[19] F. Lillo, S. Mike, and J.D. Farmer. The Long Memory of the Efficient Market. *Studies in Nonlinear Dynamics and Econometrics*, 8(3):1–35, 2004.

[20] F. Lillo, S. Mike, and J.D. Farmer. Theory for Long Memory in Supply and Demand. *Physical Review E.*, 7106:287–297, 2005.

[21] A. Lo. Reconciling Efficient Markets with Behavioural Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting*, 7(2):21–44, 2005.

[22] A.W. Lo. Long-Term Memory in Stock Market Prices. *Econometrica.*, 59:1279–1314, 1991.

[23] A.W. Lo. Reconciling Efficient Markets with Behavioral Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting.*, 7(2):21–44, 2005.

[24] A.W. Lo and A.C. MacKinlay. *A Non-Random Walk Down Wall Street.* Princeton University Press, 2001.

[25] I. Lobato and C. Velasco. Long Memory in Stock-Market Trading Volume. *Business and Economic Statistics*, 18(4):410–426, 2000.

[26] J.B. De Long, A. Shleifer, L.H. Summers, and R.J. Waldmann. Noise Trader Risk in Financial Markets. *The Journal of Political Economy.*, 98(4):703–738, 1990.

[27] K. Mannaro, M. Marchesi, and A. Setzu. Using an Artificial Financial Market for Assessing the Impact of Tobin-like Transaction Taxes. *Journal of Economic Behavior and Organization.*, 67(2):445–462, 2008.

[28] R.N. Mantegna and E.H. Stanley. Stock Market Dynamics and Turbulence: Parallel Analysis of Fluctuation Phenomena. *Physica A.*, 239(1-3):255–266, 1997.

[29] R.N. Mantegna and E.H. Stanley. Scaling Behavior in the Dynamics of an Economic Index. *Nature.*, 376:46–49, 2002.

[30] S. Martinez-Jaramillo and E. Tsang. An Heterogeneous, Endogenous and Coevolutionary GP-Based Financial Market. *IEEE Transactions on Evolutionary Computation.*, 13(1):33–53, 2009.

[31] A. Pagan. The Econometrics of Financial Markets. *Journal of Empirical Finance.*, 3(1):15–102, 1996.

[32] T. Preis, W. Paul, and J. Schneider. Fluctuation Patterns in High-frequency Financial Asset Returns. *Europhysics Letters.*, 82(6), 2008.

[33] N. Rayner, S. Phelps, and N. Constantinou. Learning Is Neither Sufficient Nor Necessary: A Dynamic Agent-Based Model of Long Memory in Financial Markets. In *Proceedings of the 13th International Conference on Electronic Commerce.*, 2011.

[34] Y. Shoham and K. Leyton-Brown. Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations. `http://www.masfoundations.org/mas.pdf`, 2010.

[35] E. Tsang, J. Li, S. Markose, H. Er, A.S Alhi, and G. Iori. EDDIE in Financial Decision Making. *Journal of Management and Economics..*, 4(4), 2000.

[36] E. Tsang, S. Markose, and H. Er. Chance Discovery in Stock Index Option and Future Arbitrage. *New Mathematics and Natural Computation, World Scientific.*, 1(3):435–447, 2005.

[37] N.J. Vriend. An Illustration of the Essential Difference Between Individual and Social Learning, and its Consequences for Computational Analyses. *Journal of Economic Dynamics and Control*, 24:1–19, 2000.

[38] F. Westerhoff. Central Bank Intervention and Feedback Traders. *Journal of International Financial Markets, Institutions and Money.*, 13(5):419–427, 2003.

[39] D.R. White and S. Poulding. A Rigorous Evaluation of Crossover and Mutation in Genetic Programming. *Proceedings of EuroGP 2009:12th European Conference on Genetic Programming*, pages 220–231, 2009.