Testing Adaptive Expectations Models of a Continuous Double Auction Market Against Empirical Facts

Neil Rayner, Steve Phelps (CCFEA CSEE) and Nick Constantinou (EBS)
University of Essex
United Kingdom

Abstract

It is well known that empirical financial time series data exhibit long memory phenomena: the behaviour of the market at various times in the past continues to exert an influence in the present. One explanation for these phenomena is that they result from a process of social learning in which poorly performing agents switch their strategy to that of other agents who appear to be more successful. We test this explanation using an agent based model and we find that the stability of the model is directly related to the dynamics of the learning process; 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 *dynamic* switching behaviour in the steady state are able to reproduce the long memory phenomena. 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.

1 Introduction

With the explosion of algorithmic trading, financial markets now constitute some of the largest and most mission critical multi-agent systems in our society. Understanding the behaviour of these markets would, it would seem, make an important contribution to the prevention of future financial crises. There is a need to build models of *actual* electronic markets, and to validate these models against empirical facts – that is, to attempt to *reverse engineer* existing multi-agent systems in order to understand how they work.

Such an exercise is now possible with the availability of electronic data detailing every transaction in the market which can run to gigabytes per year per financial asset, and can be purchased from the major financial exchanges by any third party.

Towards this end we introduce an agent based model which produces behaviour consistent with several phenomena that have been widely documented from studies of empirical financial data. Our model is in the tradition of *adaptive expectations* [Evans and Honkapohja, 2001] models in which: (i) agents' valuations are determined by their expectations of what will happen in the market in the *future*, for example their belief that the market price will rise or fall; and (ii) expectations are formed inductively through a learning process, rather than through the framework of rational expectations. This type of model is in contrast to auction theoretic models which typically assume that valuations are private information, are well defined, uncorrelated, or do not change over time, or some combination of these. In contrast the adaptive expectations framework pictures a much more dynamic view of agents' beliefs as they constantly revise their expectations, and hence valuations, in response to observations of other agents and the market itself: the market is an "expectations feedback system" from which valuations emerge [Heemeijer *et al.*, 2004].

The focus of our analysis is to determine to what extent this picture of the market is consistent with the empirical data from real exchanges. Initially we examine properties which are observable in empirical high frequency trading data with a view to model validation. Only once our model is validated can we use it answer counter factual questions such as how changes in the design of the market mechanism would affect the efficiency of the market. Therefore 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. Table 1 summarizes these phenomena.

In Section 2 we describe an existing social learning model [LeBaron and Yamamoto, 2008] and introduce our extension of it. We describe our approach to the validation of the social learning model in Section 3 and present our results in Section 4. 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. Finally in Section 7 we conclude.

2 The Model

We attempt to explain the long memory phenomena in Table 1 using 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

Table 1: Long Memory Phenomena

Long Memory	Phenomena Description
Volume	Trading volume is persistent over time [Lobato and Velasco, 2000]. Over periods of time volume can be consistently high or low.
Volatility	Stock price fluctuations have positive relations [Ding *et al.*, 1993; Engle, 1982; Mantegna and Stanley, 1997; 2002; Pagan, 1996]. Periods of similar volatility are observed ("volatility clusters").
Market Order Signs	Time series of the signs (that is, buy orders have a positive sign and sell orders have a negative sign) of market orders follows a long memory process [Lillo *et al.*, 2005].
Returns	Returns do not exhibit long memory [Cont, 2001]. Similar returns do not cluster together.

the issuing company;

2. *chartists* form valuations inductively from historical price data; and
3. *noise traders* who trade based on the fluctuations of the price of a stock. The buying and selling behaviour of traders for a particular stock generates characteristic fluctuations in price. A stock has an emergent volatility, understanding this volatility allows traders to identify when the price is relatively low and when it is relatively high.

Although chartist strategies should not be profitable according to the efficient markets hypothesis, 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. Each agent submits a limit order to the market on every round of trading. Orders are executing 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.

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 expected return $\hat{r}_{(i,t,t+\tau)}$ for the period $t+\tau$. The price of the order is set according to:

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

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

We adopt the framework of [LeBaron and Yamamoto, 2008] 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:

$$\begin{aligned}
\hat{r}_{(i,t,t+\tau)} &= \hat{r}_{f(i,t,t+\tau)} + \hat{r}_{c(i,t,t+\tau)} + \hat{r}_{n(i,t,t+\tau)} && (1)\\
\hat{r}_{f(i,t,t+\tau)} &= f_{(i,t)} \cdot (\frac{F - p_t}{p_t}) && (2)\\
\hat{r}_{c(i,t,t+\tau)} &= c_{(i,t)} \cdot \hat{r}_{L_i} && (3)\\
\hat{r}_{n(i,t,t+\tau)} &= n_{(i,t)} \cdot \epsilon_{(i,t)} && (4)
\end{aligned}$$

where F is the so-called "fundamental price" (which is exogenous and fixed for all agents), p_t is the current market quoted price which is 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 market returns over the period L_i:

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

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.

2.1 Learning

As in [LeBaron and Yamamoto, 2008], 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 over a period of 5000 units of time. Each unit of time corresponds to the entry of an order into the market by an agent. Each agent presents 5 orders to the market in this period (the number of agents in these models being 1000).

$$s_i = \frac{1}{\sum_{t=1}^{5000} (p_t - p_{(i,t+\tau)})^2} \tag{5}$$

After 5000 units of time each combination of weights held by the agents is assigned a relative fitness score (S_i) normalised with respect to the population fitness.

$$S_i = \frac{s_i}{\sum_i s_i} \tag{6}$$

The strategy weights are copied by the learning agents in proportion to this score.

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:

$$f_{(i,0)} \sim |N(0,\sigma_f)|,\ c_{(i,0)} \sim N(0,\sigma_c),\ n_{(i,0)} \sim |N(0,\sigma_n)| \quad (7)$$

In addition to the learning of weights after each 5000 units of time, agents also may mutate one of their weights drawing a weight at random from the distributions in 7.

We analyse two variants of this basic model:

1. an existing model in the literature [LeBaron and Yamamoto, 2008] in which forecasting strategies are linear *combinations* as per Equation 1 (henceforth we refer to this model as the LY model); and
2. our own model in which each agent adopts either an *atomic* fundamentalist (Equation 2) , chartist (Equation 3) or noise trader (Equation 4) forecasting rule and not a linear combination as in the LY model above (Equation 1). Both this model and the LY model occupy the same strategy space, in this model however, two out of the three weights are zero reducing each agent to just one of the return forecast rules (Equation 2, Equation 3 or Equation 4). We extend this strategy space by adding a boolean parameter which indicates if the agent strategy is contrarian or not. We use two contrarian strategies one is to negate the learnt trend, so the agent predicts a price move in the opposite direction and the second to zero the trend predicting that the price will not trend in the learnt direction but will remain at its current level. We implement this firstly by setting the contrarian chartist strategy to the negative of the non-contrarian chartist strategy:

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

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

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

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 7 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”.

3 Methodology and Model validation

We compare model assumptions according to how well a particular model reproduces the long-memory phenomena in Table 1. To compare models we test their long-memory properties using Lo’s modified rescaled range (R/S) statistic [Lo, 1991] (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-normalized 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.

Our first experiment tests the conjecture that social learning (in the LY Model) is sufficient to produce the long memory phenomena in Table 1 by attempting to reject the null hypothesis that long memory is caused by the strategies that each agent adopts and not learning at all. The model is simulated in two sequential phases with different treatment factors:

1. a learning phase in which the agent’s genetic algorithm searches for strategies with high relative fitness (see Equation 5 and Equation 6).
2. a commitment phase where agent’s commit to a learned strategy and perform no further learning.

The default experiment time is 250000 time units (taken from [LeBaron and Yamamoto, 2008]). The experiment has been executed for twice the default experiment time (2 x 250000 units of time); the learning phase is executed for half the experiment time (the default time) and then the commitment phase is started and runs for the same period (default parameter values are displayed in Table 2).

4 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 the agents move initially very quickly to a region in the strategy space. There is then a period of mean fluctuation as the agents move about in that region (not converging to any specific strategy). When the commitment phase starts and the agents stick with the strategy they have found, movement in the weights cease and we end up with a straight line for the mean value of the weights over the remainder of the experiment time with no change in the weight distribution standard deviations.

Results are shown in tables which present the percentage of executions exhibiting long memory in volume, volatility, signs of market orders (buy or sell orders) and returns for each experiment.

The experiment was executed 100 times; the results are summarised in Tables 3 and 4. In the first phase (the learning phase presented in Table 3) we see the long memory characteristics we are expecting with this model. In the second

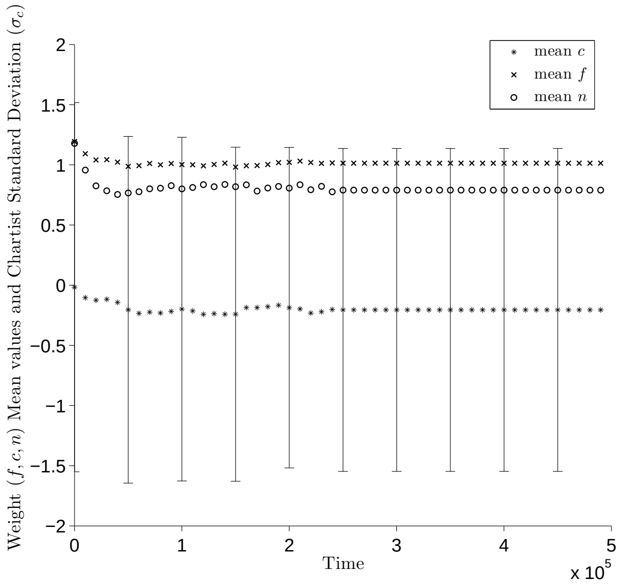

Figure 1: LY Model. Fundamentalist (f), Chartist (c), Noise (n) mean weights and Chartist Standard Deviation (σ_c) with time.

Table 2: Default Values for All Models

Parameter	Value
Std dev of fundamental weight (σ_f)	1.5
Std dev of chartist weight (σ_c)	1.5
Std dev of noise weight (σ_n)	1.5
Mutation Constant	0.08

Table 3: LY Model Learning Phase

Lag	Volume	Volatility	Market Order Signs	Returns
q=4	100	100	87	4
q=6	100	100	89	4
q=8	100	100	89	4
q=10	100	100	89	4

phase (the commitment phase presented in Table 4) we fail to generate any long memory properties.

As soon as we switch off learning these long memory phenomena disappear. It is not sufficient to have just the correct mix of strategies in order to generate long memory. So there is something about the dynamics of weight changing (caused in this case by the learning process) which is causing these phenomena. We do not have a mathematical explanation for these long memory features and yet we are able to observe that these properties emerge through the changes of agent strategy within the market.

Table 4: LY Model Commitment

Lag	Volume	Volatility	Market Order Signs	Returns
q=4	0	0	0	0
q=6	0	0	0	0
q=8	0	0	0	0
q=10	0	0	0	0

Table 5: Ranges of Parameter Values

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
Mutation Constant	0.05 to 0.20

5 Model Stability

In this section we review the stability of the LY Model. We vary some of the free-parameters described earlier, the standard deviations of the Gaussian distributions from which the weights are chosen (Equation 7) and the mutation degree. The mutation degree is the probability that any individual will mutate it's strategy and draw a new weight from the distributions in (Equation 7). We have extended the experiment execution time to highlight any problems in stability with respect to time. The experiments were run for 10 times the default time (10 x 250000 units of time) and parameter values were randomly drawn (uniformly) from the ranges in Table 5. Fifty sets of random parameter variations were executed with 10 executions for each set (totalling 500 for 2500000 units of time).

6 Model Stability Results

In Table 6 we present the results of the first experiment. We note we get negative results from the LY Model which produces weak Market Order Sign long memory but also long memory in returns (not a stylised fact of financial markets).

In Tables 7 and 8 we have separated out the long memory properties of the execution of the LY Model into an early part of the test and a later part. We note that the long memory properties of the model are changing with respect to time. The LY Model is not behaving in a stable fashion. In Table 9 we display the results for the execution of the LY Model with just the atomic extensions and finally in Table 10 we display the results for the Contrarian model. Comparing Tables 9 and 10 with the LY Model execution in Table 6 we see a substantial improvement in the stability of the Contrarian model over the LY model.

Table 6: Parameter Variation Experimental Results for LY Model

Lag	Volume	Volatility	Market Order Signs	Returns
q=4	76	86	18	12
q=6	76	83	18	14
q=8	76	82	18	14
q=10	75	81	18	15

Table 7: Early Phase Execution Results for LY Model

Lag	Volume	Volatility	Market 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 8: Later Phase Execution Results for LY Model

Lag	Volume	Volatility	Market 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

Table 9: Experimental Results for Atomic Model

Lag	Volume	Volatility	Market Order Signs	Returns
q=4	96	90	43	0
q=6	96	89	41	0
q=8	96	89	40	0
q=10	96	89	39	1

In Figure 2 the weight means of LY Model change relatively smoothly, the Contrarian Model Figure 3 in contrast, is very much more dynamic, the mean values for fundamentalist and chartist are moving a great deal relative to the LY model. Comparing the bars (which indicate the fundamentalist SD (σ_f) and chartist SD (σ_c) with time) in Figure 2 and 3 we see that the LY Model fundamentalist and chartist weight distributions tend to converge while with the Contrarian Model the chartist weight distribution hardly converges at all and the fundamentalist distribution is diverging.

Table 10: Experimental Results for Contrarian Model

Lag	Volume	Volatility	Market Order Signs	Returns
q=4	94	93	58	0
q=6	94	92	58	0
q=8	93	91	58	0
q=10	92	91	58	0

We are seeing the convergence of the LY Model into a region in the strategy space. The GA in the LY model has been successful in finding a region in this space (the successful completion of its learning). The Contrarian Model is failing to converge in the strategy space. The LY Model is not generating stable long memory properties, because the successful genetic algorithm is converging to a region in the strategy space. By restructuring the agent strategy space and increasing contrarianism we are able to retain the dynamic necessary to produce stable long memory results (Tables 9 and 10).

7 Conclusion

While imitation may contribute to the generation of long memory phenomena in real financial markets other factors must play a role in producing stable long memory phenomena over time. Our model which incorporates contrarianism and strong disparity between strategies is able to generate a more dynamic behaviour in steady state, and is therefore able to produce more realistic (stable) results. The LY model [LeBaron and Yamamoto, 2008] is not stable with respect to variation in free-parameter settings and model execution time. This was 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 modify the LY model adding atomic agents and increased contrarianism (the Contrarian Model) which retains the dynamic necessary to generate stable long memory phenomena (Section 6).

We conjecture that any model that causes and maintains a dynamic switching behaviour will produce stable long memory and that a non-learning heuristic model that mimics imitation (herding) and contrarianism would also produce positive long memory phenomena (the subject of future research).

References

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

[Ding *et al.*, 1993] 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.

[Engle, 1982] R. F. Engle. Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, 50:987–1008, 1982.

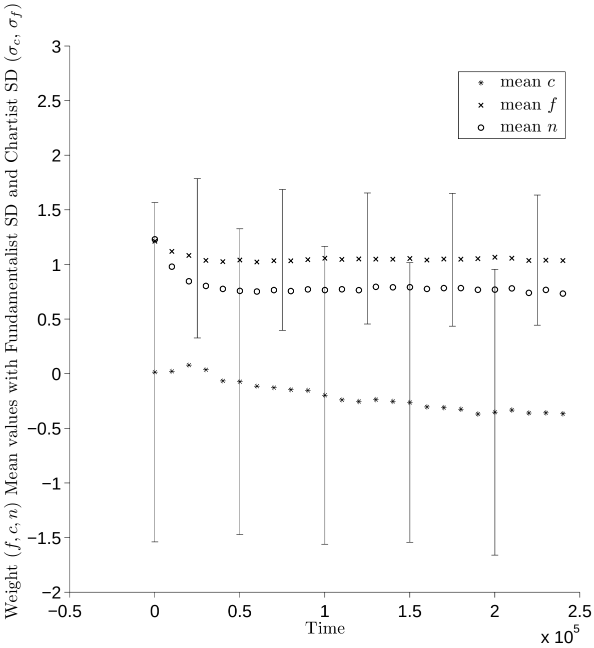

Figure 2: LY Model. Mean Fundamentalist (f), Chartist (c) and Noise (n) weights with time and Fundamentalist and Chartist Standard Deviation (σ_f, σ_c)

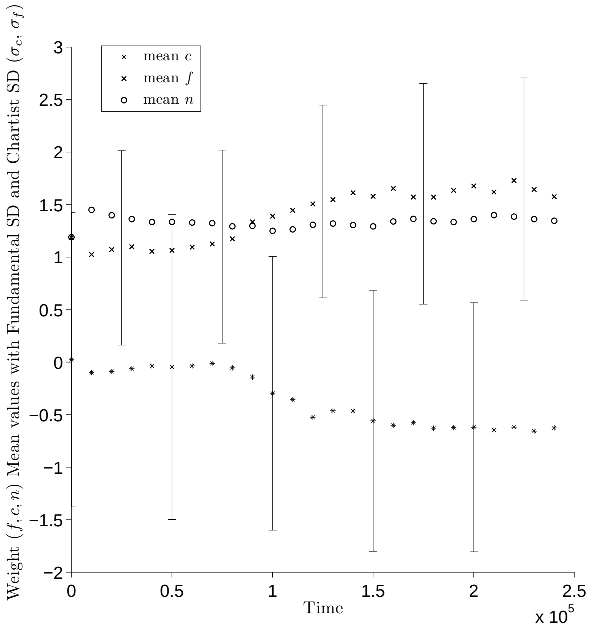

Figure 3: Contrarian Model: Mean Fundamentalist (f), Chartist (c) and Noise (n) weights with time and Fundamentalist and Chartist Standard Deviation (σ_f, σ_c)

[Evans and Honkapohja, 2001] G. W. Evans and S. Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.

[Heemeijer *et al.*, 2004] P. Heemeijer, C. H. Hommes, J. Sonnemans, and J. Tuinstra. Forming Price Expectations in Positive and Negative Feedback Systems. Technical Report 04-15, Universiteit van Amsterdam, 2004.

[LeBaron and Yamamoto, 2007] B. LeBaron and R. Yamamoto. Long-memory in an Order-Driven Market. *Physica A*, 34(1):85–89, 2007.

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

[Lillo *et al.*, 2005] F. Lillo, S. Mike, and J. D. Farmer. Theory for Long Memory in Supply and Demand. *Physical Review*, 71(6):287–297, 2005.

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

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

[Mantegna and Stanley, 1997] 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.

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

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