Learning leads to scaling behaviour in an adaptive expectations model of a double-auction market

Steve Phelps
Centre for Computational Finance & Economic Agents (CCFEA)
University of Essex, UK.
sphelps@essex.ac.uk

Wing Lon Ng
Centre for Computational Finance & Economic Agents (CCFEA)
University of Essex, UK.
wlng@essex.ac.uk

ABSTRACT

The field of econophysics has established that empirical financial time series data exhibit several robust scaling laws, but to date there has been relatively little attempt to *explain* these scaling phenomena. In this paper we explore the scaling of the absolute changes in logarithmic price with respect to the size of the time interval over which they are calculated. In an efficient market, it is straightforward to show how this can arise in a rational expectations model when changes to the fundamental price are driven by a random walk. However it is now well-known that empirical financial time series data do not fit well with geometric Brownian motion at high-frequency time scales. To address this, we model the financial market using a class of agent-based models in which agents' expectations are driven by *heuristic* forecasting rules (in contrast to the rational expectations models used in traditional theories of financial markets). We show that within this framework, social learning results in robust scaling behaviour.

1. INTRODUCTION

With the advent of algorithmic trading financial exchanges have become some of the largest and most mission critical multi-agent systems in existence. However, the recent financial crisis highlights the limitations of relying solely on theoretical models to understand these systems without validating them thoroughly against actual empirical behaviour. It is now acknowledged that widely adopted theoretical models, such as the random walk model of geometric Brownian motion, are not consistent with the data from real-world financial exchanges [13].

This had led to a resurgent interest in alternatives to models based on rational expectations models and the efficient markets hypothesis; [12] proposes the "adaptive markets hypothesis" as an alternative paradigm. The adaptive markets hypothesis posits that incremental learning processes may be able to explain phenomena that cannot be explained if we assume that agents instantaneously adopt a rational solution, and is inspired by models such as the "El Farol Bar Problem" [2] in which it is not clear that a rational expectations solution is coherent.

Appears in:

The move to electronic trading in today's markets has provided researchers with a vast quantity of data which can be used to study the behaviour of real-world systems comprised of heterogeneous autonomous agents interacting with each other, and thus a recent area of research within the multi-agent systems community [18, 18, 17, 3] attempts to take to a reverse-engineering approach in which we build agent-based models of markets that are able to replicate the statistical properties that are universally observed in real-world data sets across different markets and periods of time — the so called "stylized facts" [4].

A key issue for research in this area is that the way that agents interact and learn can be critical in explaining certain phenomena. For example, [11] introduce a model of financial markets which demonstrates that certain long-memory characteristics of financial time series data can only be replicated when agents imitate each others' strategies. When their model is analysed under a treatment in which learning does not occur, the corresponding long-memory properties disappear. [12] posits that many features observed in financial time series data are the result of adaptation; that is, heuristic learning methods that iteratively refine solutions.

In this paper, we use a similar model to explore whether adaptation could be responsible for some of the *scaling properties* observed in empirical financial time-series data, which we describe below.

[15] coined the term coastline paradox and used Britain as an example to explain that the measured length of the coastline is highly dependent on the length of the yardstick that was used for the assessment. A shorter yardstick would encompass more of the tinier meanders along the coast, whereas a longer one would simply ignore the finer sinuosities; therefore, the estimated coastline will increase, possibly to infinity, as the length of the yardstick decreases.

A similar "fractal" behaviour can be found in economic or financial times series such as price curves: the "coastline" of a price curve is the result of all price movements. At the highest resolution, i.e. the shortest possible interval between two observations, price changes appear miniscule but they ultimately form the entire price coastline and, being the source of price variation, they determine its length. As [6] pointed out, the sum of all price changes for a given time interval of measuring can be far longer than one would have expected when comparing it to low resolution levels.

The main advantage of scaling laws is their universality and scale invariance, allowing for both flexibility and consistency in modelling. Scaling laws help to detect whether observed phenomena are to a certain degree similar or even the same at different scales. The notable feature of this self-similarity concept is that the characteristics and their implications would apply to both short-term and longer-term behaviour of price dynamics. [6], for instance, illustrated on high-frequency foreign exchange data how one one can exploit the plethora of intra-day data to construct robust trading models and then simply scale models to address both short-term shocks and long-term fluctuations in market movements, benefitting from the scale invariance property of the scaling law.

The outline of this paper is as follows. In the following section we describe the agent-based model used for our analysis. In Section 3 we describe how we test for the presence of scaling behaviour in this model. We present our findings in Section 4. Finally, we conclude in Section 5.

2. THE MODEL

We use a class of agent-based model that has been demonstrated in previous studies by other researchers to replicate many of the statistical properties observed in empirical data [10, 11, 18][1]. The model simulates not only the detailed micro-structural operation of a financial market, but also how agents form their valuations for the asset being traded based on observations of past prices and the bidding behaviour of others; that is, using the terminology of auction-theory, we model the financial market as an interdependent values scenario.

The market is modelled as an order-driven exchange, typical of that used to trade equities in electronic markets such as the London Stock Exchange (LSE). In an order-driven exchange, agents can submit limit-orders which are offers to buy, or sell, at a specified price and quantity. Orders are matched using a continuous double auction (CDA) [5]. A buy order can be matched with a sell order if the buy price is greater than the sell price, and vice versa. When new orders arrive they are executed immediately at the price of the earliest order, provided that a corresponding match can be found. The fulfilled orders are then removed from the exchange. If orders cannot be matched immediately, they are queued on the "order book" until either they are matched or their expiry time is reached. Buy orders and sell orders each have their own priority queue and are ranked in descending and ascending order respectively. The highest outstanding buy order and the lowest sell order are called the best bid and the best ask respectively, and the pair of prices corresponding to the best bid and ask is called the market quote.

The expectations model is based on an "adaptive expectations" framework [12] in which agents make trading decisions based on forecasts of the next period price, which are then updated through a social learning process. Agent's expectations are modelled according to a chartist, fundamentalist and noise framework [14]. Chartists believe future prices can be forecast by extrapolating from past prices, in contradiction to the efficient markets hypothesis. On the other hand, fundamentalists believe that there is a fair price for the asset, to which the future price will revert. In a market in which not every agent is rational, the best class of rule to use is not necessarily the fundamentalist forecast. For example, if a significant fraction of the market uses chartist rules, this may create trends away from the fair price, and an agent using a fundamentalist rule might do better by switching to a chartist rule; i.e., chartist expectations can become self-fulfilling. Finally, noise traders form their expectations independently at random and trade on this noise believing that it is a signal.

The model is implemented as a discrete-event simulation of an entire trading day. The trading day is divided into 2×10^5 discrete time steps, each of equal duration. This gives a resolution of approximately 150 milliseconds on a real-time scale. The unequal spacing of events that are typically observed in empirical market data sampled at high-frequency is modelled using a Bernoulli process; at each time step, an randomly chosen agent arrives at the simulation with probability λ. Inter-arrival times are therefore approximately exponentially distributed, as per a continuous Poisson arrival model.

Each agent maintains a single position in the market which it revises according to its valuation policy. The market is populated with a total of n agents. When the i^{th} agent arrives at the simulation it either places a new limit order, or amends its existing order, based on its valuation $v_{(i,t)}$. The price of the order is calculated using a strategy similar to the Zero-Intelligence Constrained (ZI-C) behaviour described in [7]; every agent maintains a randomly chosen markup $\delta_i \sim U(0, \delta_{max})$ and the price of the order is set to

$$\rho_{(i,t)} = v_{(i,t)}(1 + \delta_i) \quad (1)$$

for a sell order, or

$$\rho_{(i,t)} = v_{(i,t)}(1 - \delta_i) \quad (2)$$

for a buy order. The direction of the order (buy or sell) is determined by the agent's expectation of the next period price v_i compared to the current market price p_t: if $v_{(i,t)} > p_t$ then a buy order is placed (the agent takes a long position), otherwise a sell order (short position).

Agents decide their valuations $v_{(i,t)}$ as a function of the financial returns of the asset, i.e. the time-series of changes in the logarithmic prices sampled at a particular frequency:

$$r_{j,\Delta_t} = log(p_{j\Delta t}) - log(p_{(j-1)\Delta_t}) \quad (3)$$

where p_t is the market price observed at time t, and Δ_t is the sampling interval. The market price p_t is defined as either: the price of the transaction that occurred at time t, or the middle of the quote if no transaction occurred.

2.1 Valuations

The valuation $v_{(i,t)}$ of the i^{th} agent is calculated by making a forecast of the next period logarithmic return $\hat{r}_{(t,i)}$. Agents using a chartist policy ($\hat{\text{rc}}$) use a simple moving average of past returns:

$$\hat{\text{rc}}_{(i,t)} = \sum_{j=t/\Delta t - w_i}^{t/\Delta t} r_{j,\Delta_t}/w_i \quad (4)$$

where w_i is the window size used by agent i. Agents using a

9[1]Although this is not *our* model, we devote space to describing it here, as well as in one of our other submissions [1], in order to make the paper self-contained, and also because there are some subtle technical differences between our implementation and existing models. Firstly, chartist return forecasts are calculated using logarithmic returns rather than simple returns in order to avoid a bias towards positive drift. Secondly, learning occurs probabilistically at every time step rather than being scheduled at every 5000 steps in order to avoid artifacts relating to strong temporal synchronisation of learning. Finally, in line with the learning-classifier system literature, we use an exponential moving average of forecast errors in order to determine fitness rather than a moving window.

fundamentalist forecasting rule (rf) have information about the fair value of the asset f_t and forecast accordingly:

$$\hat{\mathrm{rf}}_{(i,t)} = log(p_t) - log(f_t). \quad (5)$$

Agents using a noise-trader rule (rn) make random return forecasts:

$$\hat{\mathrm{rn}}_{(i,t)} = \epsilon_t \quad (6)$$

where ϵ_t are *i.i.d.* random variates drawn from a standard normal distribution $N(0,1)$.

As in [10], we allow agents to make forecasts using a linear combination of all three types of rule in order to estimate the return over the next time period τ_i. Let

$$\boldsymbol{R}_{(i,t)} = (\hat{\mathrm{rf}}_{(i,t)}, \hat{\mathrm{rc}}_{(i,t)}, \hat{\mathrm{rn}}_{(i,t)}) \quad (7)$$

denote the vector of return forecasts made according to each of the fundamentalist, chartist and noise rules respectively. Let $\boldsymbol{S}_{(i,t)} \in \mathbb{R}^3$ denote a vector of coefficients which define the forecasting strategy of agent i at time t. Then the return forecast $\hat{r}_{(i,t)}$ is given by:

$$\hat{r}_{(i,t)} = \boldsymbol{S}_{(i,t)} \cdot \boldsymbol{R}_{(i,t)}. \quad (8)$$

The valuation of agent i is then given by their forecast of the price at time $t+\tau$:

$$v_{(i,t)} = p_t e^{\hat{r}_{(i,t)}/\tau_i}. \quad (9)$$

This can then be substituted into equations 1 and 2 in order to determine the agent's limit price.

In order to model the arrival of market news, changes in the growth rate of the fundamental price follow a random walk:

$$df_t = \mu_f f_t + \sigma_f f_t \; dW_t \quad (10)$$

where σ_f is the volatility, μ_f is the drift and dW_t is a standard Wiener process discretised over the time interval Δt:

$$dW_t \sim \sqrt{\Delta t} N(0,1) \quad (11)$$

2.2 Initial conditions

The initial values for agents' strategy coefficients are random variables $S_{(i,0)} = (\mathrm{SF}, \mathrm{SC}, \mathrm{SN})$ with distributions:

$$\begin{aligned} \mathrm{SF} &\sim |N(0,\sigma_f)|, \\ \mathrm{CF} &\sim N(0,\sigma_c), \\ \mathrm{SN} &\sim |N(0,\sigma_n)| \end{aligned} \quad (12)$$

where σ_f, σ_c and σ_n are the standard deviations of the fundamentalist, chartist and noise components respectively. In a static experiment treatment these values remain constant over the course of a simulation run. Under a learning treatment, they evolve over time according to the learning model specified below.

2.3 Learning

As in [11, 18], we model adaptive expectations [12] by allowing agents to learn their forecasting strategy $\boldsymbol{S}$ using a model of social learning implemented in the form of a simple co-evolutionary genetic algorithm.

Each agent records the exponential moving average of its forecast error as the market progresses:

$$\bar{\eta}_{(i,t)} = \alpha(\hat{r}_{(i,t)} - r_{(i,t)})^2 + (1-\alpha)\bar{\eta}_{(i,t-w_\eta)}. \quad (13)$$

Symbol	Definition
$n \sim U(80,120)$	number of traders
$\sigma_c \sim U(0.1,3)$	Standard deviation of chartist distribution
$\sigma_n \sim U(0.1,3)$	Standard deviation of noise-trader distribution
$\sigma_f \sim U(0.1,3)$	Standard deviation of fundamentalist distribution
$\delta_{max} \sim U(0,1)$	Maximum value of the markup distribution
$\lambda \sim U(0.1,0.9)$	Probability of agent arrival per time step
$\lambda_r \sim U(0.1,0.9)$	Imitation probability per agent per time step
$\lambda_m \sim U(0.1,0.2)$	Mutation probability per agent per time step
$w_i \sim U(1,100)$	Chartist window size
$w_\eta \sim U(1,100)$	Sampling interval for forecast errors
$\tau_i \sim U(5,10)$	Forecast time horizon
$\sigma_f \sim U(0,1.0)$	Volatility of fundamental price process
$\mu_f \sim U(0,0.1)$	Drift of fundamental price process
$p_0 \sim U(100,100)$	Initial price

Table 1: Summary of model parameters and their associated distributions

The fitness of agent i is then given by

$$\phi_{(i,t)} = 1/(1+\bar{\eta}_{(i,t)}). \quad (14)$$

Imitation occurs with probability λ_r per time step for any given agent. When imitation occurs, i imitates another agent by randomly selecting a partner j from the remainder of the population with probability proportionate to fitness:

$$p(agent_{j \neq i}) = \phi_{(j,t)} / \sum_{k \neq i} \phi_{(k,t)}, \quad (15)$$

and then agent i inherits its strategy from agent j; that is, $\boldsymbol{S}(i,t) = \boldsymbol{S}_{(j,t-1)}$.

Mutation occurs with probability λ_m per time step for any given agent. When mutation occurs the agent re-initialises its strategy by redrawing its coefficients $\boldsymbol{S}$ from the distributions specified in Section 2.2.

3. METHODOLOGY

We analyse the model specified the previous section using Monte-Carlo simulation under two different treatment conditions: *learning* versus *static*. Under the former treatment, agents' strategies remain fixed throughout the duration of the trading day, whereas under the learning treatment the strategy of each agent S_i evolves according to the social-learning model specified in Section 2.3. For each treatment we run total of 200 independent simulations with free parameters drawn *i.i.d.* from the distributions specified in Table 3. We then analyse the scaling properties of the resulting price time-series as described in the following section.

3.1 Scaling Laws

In its simplest form, a scaling law (or sometimes also referred to as power law) is a polynomial relationship

$$f(x) = ax^b \tag{16}$$

that satisfies the property of scale invariance

$$f(cx) \propto f(x) \tag{17}$$

for $a, b, c \in \mathbb{R}$. A rescaling of the argument x only changes the factor of proportionality but not the functional relationship itself

$$f(cx) = a(cx)^b = c^b ax^b = c^b f(x) \propto f(x). \tag{18}$$

Likewise, *changes* of x and $f(x)$ are also proportional; this becomes clearer when looking at their logarithms

$$\log(f(x)) = log(a) + b\log(x). \tag{19}$$

This scale invariance of $f(x)$ is a the main feature of scaling laws, allowing the researcher to model complex phenomena in a very efficient manner [6]. For example, consider a standard Random Walk

$$X_t \quad = \quad X_{t-1} + \varepsilon_t = \sum_{i=0}^{t} \varepsilon_i \tag{20}$$

with $X_0 = 0$ and Gaussian White Noise innovations $\varepsilon_t \sim \mathcal{N}(\mu_\varepsilon, \sigma_\varepsilon^2)$. As the innovations ε_t are i.i.d., the Random Walk's mean and variance over a longer horizon, here denoted as $\mu_X(t)$ and $\sigma_X^2(t)$, can be straightforwardly calculated as a multiple of the innovations' moments

$$\mu_X(t) \quad = \quad E\left(\sum_{i=0}^{t} \varepsilon_i\right) = \sum_{i=0}^{t} E\left(\varepsilon_i\right) = t\mu_\varepsilon \tag{21}$$

$$\sigma_X^2(t) \quad = \quad Var\left(\sum_{i=0}^{t} \varepsilon_i\right) = \sum_{i=0}^{t} Var\left(\varepsilon_i\right) = t\sigma_\varepsilon^2. \tag{22}$$

This so-called Gaussian scaling law that describes the first two moments of the Random Walk at time t simply as a function of multiples of the White Noise's moments is often applied in finance, for example, to model volatility by rewriting eq. (22) as

$$\sigma_X(t) = \sigma_\varepsilon \cdot t^{1/2} \tag{23}$$

(similar to eq. (16) with $a = \sigma_\varepsilon$ and $b = 1/2$).

Many empirical studies have reported that financial asset returns do not have Gaussian distributions but exhibit "stylised facts" [4]. It is well-known that empirical return innovations are not i.i.d., therefore point forecasts for the volatility will be biased or noisy. However, the forecast for the *cumulative* volatility (22) will become more accurate because of errors cancellation and volatility mean reversion, allowing for more robust estimates [16].

Consider the return definition in eq. (3). One popular scaling law models the size of the average absolute price change $E\left(|r_{\Delta_t}|\right)$ as a linear function of the time interval Δt of its occurrence

$$E\left(|r_{\Delta_t}|\right) \propto \Delta t^b, \tag{24}$$

where b refers to the Hurst exponent that is often estimated in long range dependence analysis. Eq. (24) is particularly interesting for applied risk management as it suggests that that the expected absolute price change is proportional to the elapsed time interval raised to a power b.

Consider eqs. (16) and (24). Now define the absolute return as

$$X_{\Delta t} = |r_{\Delta_t}| \tag{25}$$

and its mean as

$$E(X_{\Delta t}) = \langle X_{\Delta t} \rangle \,. \tag{26}$$

Note that $\langle X_{\Delta t} \rangle$ is a random variable itself and depends on the sampling interval Δt (see also eqs. (21) and (22)).

The interval Δt is a continuous variable. To obtain a range of suitable discrete sampling intervals, which we denote as ι, we first arbitrarily choose a base of 1.2 and consider the series of its powers in increasing order, i.e. 1.2^1, $1.2^{1.5}$, 1.2^2, ... Calibrating against empirical transaction data obtained from the London Stock Exchange to approximate a 5 minute interval, we determine that the lowest appropriate value of Δt is $\iota_1 = 342 = \lceil 1.2^{32} \rceil$. The subsequent values are simply higher powers (rounded) of the same base 1.2, i.e. $\iota_2 = 374 \approx 1.2^{32.5}, \iota_3 = 410 \approx 1.2^{33}, \iota_4 = 449 \approx 1.2^{33.5}, \ldots$, yielding $\iota = \{342, 374, 410, \ldots, 1225\}$. These values represent the number of discrete time steps in the agent-based models. ι_{max} is set to 1225 to avoid obtaining too few data points which would be the case if the time gap between two observations would be even larger.

We sample the price data in these different discrete time intervals ι_i and then compute the corresponding absolute price changes and their means. Figures 1 and 2 illustrate this procedure with ι_1 and ι_8 as examples. Although ι_8 is almost twice as long as ι_1, the resulting distributions of both log-returns (Figure 1) and absolute returns (Figure 2) do not seem to differ much in their characteristics but rather similar. It is this notable feature of self-similarity that allows the researcher to analyse both short-term and longer-term price dynamics in a consistent manner, because rescaling the argument ι only changes the factor of proportionality but not the functional relationship itself.

If the relationship

$$\langle X_{\Delta t} \rangle = a(\Delta t)^b \tag{27}$$

holds, it would follow that

$$E(X_{c\Delta t}) = c^b E(X_{\Delta t}) \propto E(X_{\Delta t}), \tag{28}$$

implying that a change of Δt would also result in a *proportional* change of $\langle X_{\Delta t} \rangle$, relating price changes from both long-term and short-term time scales.

This remarkable feature of the scaling law approach would allow us to make more robust forecasts on the magnitude of expected price changes for multiple time frames, *regardless* of the chosen sample period due to the concept of self-similarity from fractal theory. The expected price change $\langle X_{\Delta t} \rangle$ for any time horizon can then be estimated by running standard OLS (Ordinary Least Squares) regressions on the logarithmic version of eq. (27)

$$\log(\langle X_{\Delta t} \rangle) = \tilde{a} + b\log(\Delta t) \tag{29}$$

with $\tilde{a} = \log(a)$ (see also eq.(19)).

Here, $\tilde{a}$ is the intercept (a is not estimated but $\log(a)$) and the slope b represents the scaling exponent (or fractal dimension). If $\tilde{a}$ is significant, which can be tested with a t-test, then $\exp(\tilde{a}) = a$ can be interpreted as the expected "minimum spread" that a trader has to pay or can receive, *regardless* of the elapsed time since the opening of his position. In contrast, b is in this log-log model measures the

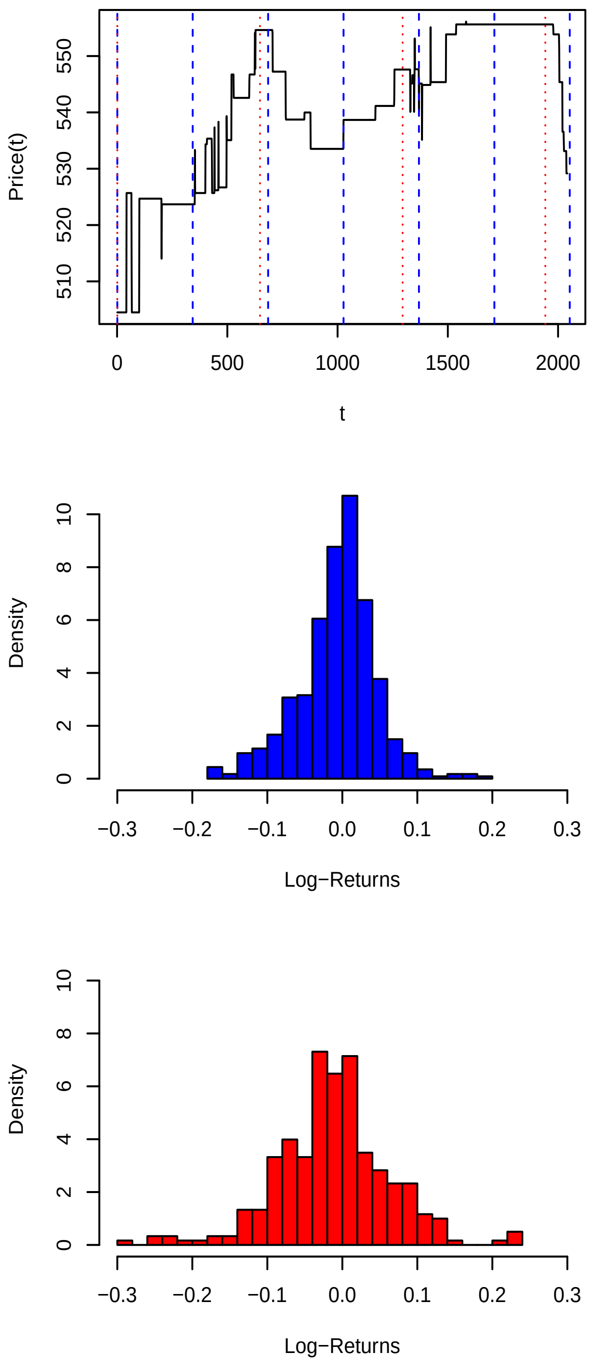

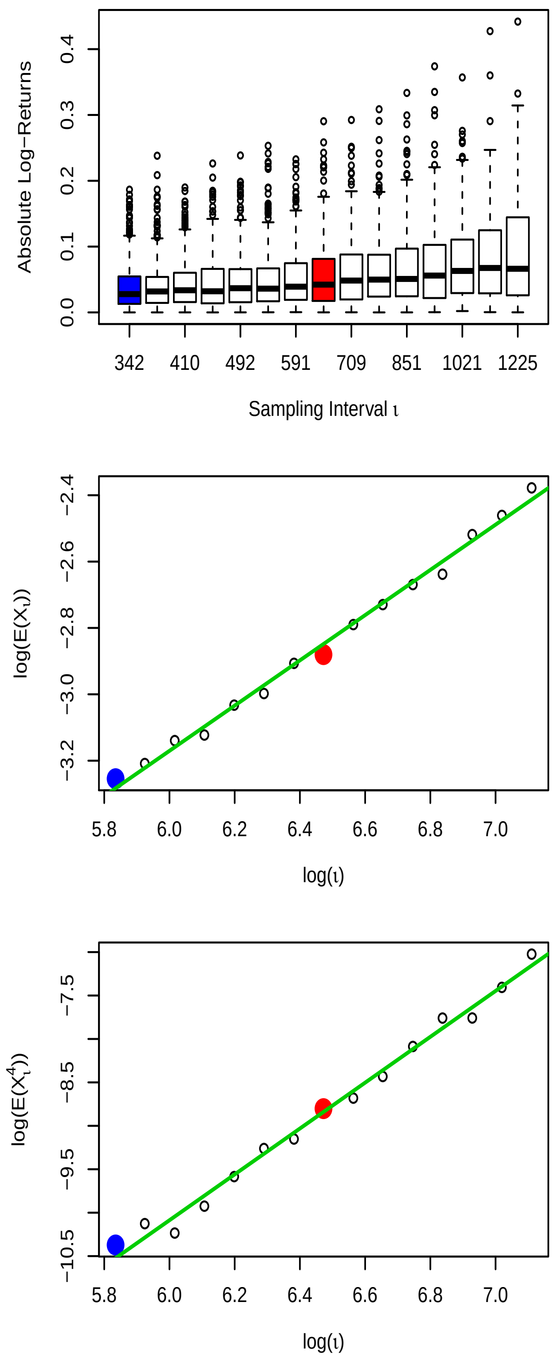

Figure 1: The top panel shows (an extract of) a randomly chosen sample price series generated by the learning model. As it can be seen, a very small sampling interval would yield many zero returns and induce possibly spurious strong autocorrelation. The blue dashed and red dotted vertical lines indicate the position of the prices sampled at frequency $\iota_1 = 342$ and $\iota_1 = 647$ steps, respectively. The distribution of the corresponding returns is illustrated in the histograms in the middle (blue) and bottom (red) panel. As expected, a higher sampling interval reduces the amount of (close to) zero returns and increases the amount of extreme returns.

Figure 2: The top panel compares boxplots of the *absolute* returns sampled at the different intervals ι_i, exhibiting their 'self-similarity' at different timescales. For example, the distributions of the absolute values of the returns shown in Figure 1 are represented here in the far left blue boxplot and in the red boxplot in the center. Now consider only the means of the absolute returns raised to the power $pow = 1$ (y-axis in middle panel) or $pow = 4$ (y-axis in bottom panel). The parameters of the corresponding scaling law in eq. (31) are then obtained by regressing these averages against the sampling intervals ι_i in the log-log dimension.

percentage change in the average absolute price change with respect to the percentage change of the considered time interval.

Once, we have calculated $\langle X_\iota \rangle$ for all ι_i (see Figure 2 top panel), we estimate the scaling law parameters according to eq. (29). The intercept and slope of the green line in the middle panel in Figure 2 correspond to parameters $\tilde{a}$ and b in eq. (29).

As mentioned above, a major advantage of scaling laws is their universality and scale invariance, allowing for both flexibility and consistency in modelling. A slightly modified version of eq. (24) that considers higher (and possibly non-integer) moments of the absolute price change can be specified as

$$(E(|r_{\Delta_t}|^{pow})) \propto \Delta t^b, \quad (30)$$

where *pow* corresponds to the power of the moment. The popular realised variance measure in the financial literature can be seen as a special of eq. (30) case with $pow = 2$, likewise eq. (23). In contrast to the financial literature which mostly only consider the special cases $pow = 1, 2$ [6], we also consider a generalisation similar to eq. (30), also to assess the robustness of the scalability of our results.

In particular, we are interested in the relationship

$$\langle X_{\Delta t}^{pow} \rangle = a_{pow}(\Delta t)^{b_{pow}}, \quad (31)$$

where *pow* refers to the power of the moment of interest, e.g. $\langle X_{\Delta t}^2 \rangle = E\left(|r_{\Delta_t}|^2\right)$ (see eq. 23). If we were interested in the kurtosis of the returns distribution, i.e. their fourth moment, to assess the their "fat tails" and extreme events, we can consider eq. (31) and estimate the corresponding parameters $\tilde{a}_4$ and b_4. As it can be seen in the bottom panel in Figure 2, a scaling law relationship can also be found for the quartic order.

This demonstrates again that the scale invariance of scaling laws (31) offers great benefits to researchers for their modelling as it allows them to describe complex phenomena in a very efficient and consistent manner without changing or making assumptions depending on the data sample period of sampling frequency. In the following, we estimate the parameter sets $(\tilde{a}_{pow}, b_{pow})$ for the scaling law (31) with $pow = \{1.0, 1.2, 1.4, \ldots, 5\}$. This procedure is then repeated for all price series in all agent based models. Our scaling law methodology is briefly summarised in Listing 1.

4. RESULTS

Our goal is to analyse the which of the agent-based model parameters listed in Section 3.1 influences the scaling behaviour of a particular price building process as described in eq. (31).

For better comparison reasons, we also include in our analysis a standard Geometric Brownian Motion (GBM), whose theoretical scaling exponent can be derived as $b_{pow} = pow/2$, as a reference model (see Listing 1). Across all agent-based models, the vast majority of the simulated time series have pronounced scaling properties (see example in Figure 2), because all most all estimated coefficients $(\widehat{\tilde{a}_{pow}^k}, \widehat{b_{pow}^k})$ (see eqs. (29) and (31)) are different from zero at the 1% significance level, and the corresponding R^2 from the regression is larger than 95%.

In the following discussion we will not comment on all 12600 individual regressions (3(models)$\times$ 200(time series)$\times$

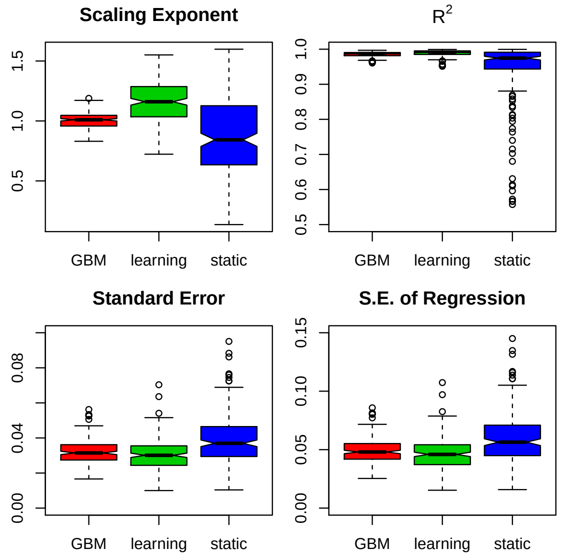

Figure 3: Comparison of the quality of the scaling law regression results for $pow = 2$ across all models. The boxplots show the distribution of estimated scaling exponents (slope parameters) b_2^k (top left panel), their corresponding standard errors (bottom left), the R^2 of the regression eq. (29) (top right), and the standard error of the regression (bottom right).

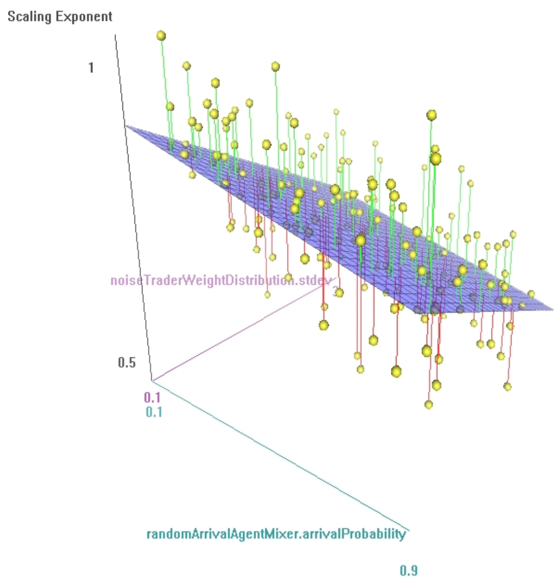

Figure 4: Illustration of the *combined* linear effect of the agent-based model parameters σ_n (standard deviation of the noise trader's weight distribution) *and* λ (probability of agent arrival per time step) on the estimated scaling exponent $\widehat{b_2^k}$ in the learning-network model.

Algorithm 1: Estimate scaling law parameters $(\tilde{a}_{pow}, b_{pow})$

Data: Log-Price Series $log(p_t)$
Result: $(\tilde{a}^k_{pow}, b^k_{pow})$ of all price series for all three data generating processes
Consider the models:
1. Geometric Brownian Motion
2. learning
3. static

begin
  Select model $m = \{1, 2, 3\}$
  for *sample price process* $k = 1$ **to** $k = 200$ **do**
    for *power* $pow_1 = 1$ **to** $pow_{max} = 5$ **do**
      for *time scale* ι_1 **to** ι_{max} **do**
        Sample p_t at $t = \{0, \iota, 2\iota, 3\iota, ...\}$
        Calculate X_ι
        Calculate $\langle X^{pow}_\iota \rangle = E\left((X_\iota)^p\right)$
        Regress $\log(\langle X^{pow}_\iota \rangle) = \tilde{a}_{pow} + b_{pow}\log(\iota)$
        Save estimates as $(\widehat{\tilde{a}^k_p}, \widehat{b^k_p})$
      end
    end
  end
end

21(moments)), but rather only focus on the special case $pow = 2$, due to its popularity in the finance literature [16]. However, all obtained results are scalable for all considered timescales and the implications apply to all other moments in a similar way.

Figure 3 compares the distribution of estimated scaling exponents (slope parameters) $\widehat{b^k_2}$ for all agent-based models. Two major findings are evident. Firstly, the notched boxplots in the top left panel show that the scaling exponent from the model involving learning is on average higher than the corresponding value from the static model. A higher scaling exponent b_2 implies that the expected squared price change increases for a given ι (see eq. (31)). Compared to both the GBM and the static models, learning causes more variation in the returns.

The second major result is that the average scaling exponent of the static model is not just lower than compared to the GMB, its range of estimated values is much wider compared the other models, although the fundamental price is following a GBM price path. This finding is supported by the fact that the corresponding R^2s are lower (Figure 3, top right panel), and that the standard error of the estimates (bottom left panel) and the standard error of the regression (bottom right panel) higher, all indicating a poorer fit in the scaling property.

Parameter	learning	static
n	6.6E-5 *** (2.760)	1.2E-5 (0.037)
σ_c	1.1E-3 ** (2.178)	7.2E-3 (0.927)
σ_n	-1.1E-3 * (−1.744)	2.4E-2 (1.160)
λ_m	-1.1E-2 (−1.449)	-2.3E-1 ** (−2.080)

Table 2: Measuring the influence of the four agent-based model parameters n (number of agents), σ_c (standard deviation of the chartist's weight distribution), σ_n (standard deviation of the noise trader's weight distribution) and λ_m (mutation probability) on the obtained R^2 of the corresponding scaling regressions. t-statistics are shown in parentheses. The significance levels are : '*' = 1%, '**' = 5%, '*'= 10%.**

It is to be emphasised that the bigger dispersion in the distribution of the estimates $\widehat{b^k_2}$ is not due to noise of induced bias but the actual mechanism in the agent-based model. Regressing $\widehat{b^k_2}$ against all agent-based model parameters (this is not to be confused with eq. (29)) reveals that particularly two decisive factors are consistently controlling the price processes' scaling behaviour in all moments. Figure 4 illustrates the *combined* effect of σ_n (the standard deviation of the noise trader's weight distribution) and λ (the probability of agent arrival per time step) on the scaling exponent $\widehat{b^k_2}$ in the learning network model. In general, both variables have a negative impact in the scaling exponent, meaning that the higher two variables are in the agent-based model's setting, the lower are the expected price variation at all timescales. These effect can be explained such that first a higher standard deviation of the noise traders' weights will create a wider dispersion of price expectations around the midprice, making the more distant prices less likely to be executed. The remaining prices closer to the midprice are then more considered much more often, yielding lower price variation in the long run. Secondly, a higher arrival rate of agents improves the market's liquidity in general, but ultimately only the more efficient prices will be accepted by the market resulting in price fluctuation of smaller magnitudes.

In order to assess which and how certain factors affect the goodness of fit in the scaling property, we also regress the obtained R^2 against theagent-based model parameters. The corresponding multiple regressions results are listed in Table 2 (for reasons of simplicity, only variables that are influencing the fit of either model are shown in this table). In the learning model both n (number of agents) and σ_c (standard deviation of the chartist's weight distribution) increases the fit of the scaling property, whereas σ_n (standard deviation of the noise trader's weight distribution) decreases it. In the static model however, only λ_m (mutation probability) has a statistically negative significant on the obtained R^2 of the corresponding scaling regressions.

5. CONCLUSION

In this paper we investigated whether the scaling characteristics observed in empirical financial time-series data could be explained by a process of social learning, in accordance with the "adaptive markets hypothesis" [13].

To this end, we used an agent-based model based on earlier models [10, 11] which have been demonstrated to reproduce many other stylized facts of financial markets including volatility-clustering, fat-tailed return distributions and long-memory properties. Previous work has shown the importance of learning in reproducing long-memory features of the price process [11]. In this paper we have shown that learning also contributes to realistic scaling behaviour. To the best of our knowledge, this is the first attempt to systematically analyse which features of of an adaptive-expectations model contribute to scaling behaviour in financial markets.

We have shown that learning results in more robust scaling behaviour, as measured by the R^2 of the scaling regression, when compared to a control condition in which learning does not occur (the static treatment).

Furthermore, we found that two distinct variables, i.e. the noise trader weight distribution's standard deviation and the probability of agent arrival per time step, have a particularly strong and statistically significant impact on the scaling exponent. Across all simulated price paths, the higher these two variables are in the model setting, the lower are the expected price variation, consistently at all timescales. Whilst we cannot directly account for these correlations under the present analysis, we note that both of these parameters play an important role in the learning process described in Section 2.3. Other research has highlighted the importance of the *dynamics* of learning [18] in reproducing particular statistical features observed in empirical data, and [9, 19, 20] demonstrate the existence of chaotic attractors in strategy dynamics in several theoretical settings including an expectations framework. Moreover, there is evidence of oscillating behaviour in the learning dynamics exhibited by human subjects in laboratory asset-pricing experiments [8]. An interesting research question is whether the scaling properties exhibited in our current study are the result of similar chaotic attractors in the co-evolutionary dynamics resulting from social learning.

The simulated agent based models provide an abundance of data. In this study, we only considered the scaling behaviour of absolute prices changes. Many other variables such as trading durations, trading volume, number of trades remain to be analysed. The investigation of these and further variables can eventually shed more light on how the dynamics in the market's the supply and demand can be modelled in a universal scale invariant ways. In our future research we will focus on further scaling laws for other financial variables to study their fractal nature and degree of self-similarity.

Acknowledgements

We would like to thank Khurom Kyani for contributing many helpful discussions and ideas.

6. REFERENCES

[1] #536. Network motifs for microeconomic analysis. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, 2013.

[2] W. B. Arthur. Inductive Reasoning and Bounded Rationality. *The American Economic Review*, 84(2):406–411, 1994.

[3] B.-a. Cassell and M. P. Wellman. Agent-Based Analysis of Asset Pricing under Ambiguous Information: an empirical game-theoretic analysis. *Computational & Mathematical Organization Theory*, 2012.

[4] R. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

[5] D. Friedman and J. Rust, editors. *The Double Auction Market: Institutions, Theories, and Evidence (Proceedings of the workshop on double auction markets held June 1991 in Santa Fe, New Mexico)*. Westview, 1993.

[6] J. B. Glattfelder, A. Dupuis, and R. B. Olsen. Patterns in high-frequency FX data: discovery of 12 empirical scaling laws. *Quantitative Finance*, 11(4):599–614, 2011.

[7] D. K. Gode and S. Sunder. Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137, 1993.

[8] C. Hommes, J. Sonnemans, J. Tuinstra, and H. van de Velden. Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133, 2008.

[9] C. Hommes and G. Sorger. Consistent Expectations Equilibria. *Macroeconomic Dynamics*, 2(3):287–321, Sept. 1998.

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

[11] B. LeBaron and R. Yamamoto. Long-memory in an order-driven market. *Physica A: Statistical Mechanics and its Applications*, 383(1):85–89, Sept. 2007.

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

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

[14] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(February):498–500, 1999.

[15] B. B. Mandelbrot. How long is the coast of Britain? Statistical self-similarity and fractional dimension. *Science*, 156(3775):636–638, 1967.

[16] M. McAleer and M. Medeiros. Realized Volatility: A Review. *Econometric Reviews*, 27(1):10–45, 2008.

[17] I. Palit, S. Phelps, and W. L. Ng. Can a Zero-Intelligence Plus Model Explain the Stylized Facts of Financial Time Series Data? In *Proceedings of the Eleventh International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) - Volume 2*, pages 653–660, Valencia, Spain, 2012. International Foundation for Autonomous Agents and Multiagent Systems.

[18] N. Rayner, S. Phelps, and N. Constantinou. Testing adaptive expectations models of a double auction market against empirical facts. In *Lecture Notes on Business Information Processing: Agent-Mediated Electronic Commerce and Trading Agent Design*, Barcelona, 2011. CCFEA, Springer (in press).

[19] Y. Sato, E. Akiyama, and J. D. Farmer. Chaos in learning a simple two-person game. *Proceedings of the National Academy of Sciences of the United States of America*, 99(7):4748–4751, Apr. 2002.

[20] B. Skyrms. Chaos in Game Dynamics. *Journal of Logic, Language, and Information*, 1:111–130, 1992.