King's Research Portal

DOI:

[10.1002/isaf.1346](https://doi.org/10.1002/isaf.1346)

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Phelps, S., & Ng, W. L. (2014). A simulation analysis of herding and unifractal scaling behaviour. *Intelligent Systems in Accounting, Finance and Management*, 21(1), 39-58. <https://doi.org/10.1002/isaf.1346>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

A simulation analysis of herding and unifractal scaling behaviour

Steve Phelps* Wing Lon Ng†

Centre for Computational Finance and Economic Agents (CCFEA)

University of Essex, Wivenhoe Park, Colchester CO4 3SQ, UK

18th May 2017

Abstract

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, we can reproduce unifractal scaling with respect to three well-known power-laws relating: (i) moments of the absolute price-change to the time-scale over which they are measured; (ii) magnitude of returns with respect to their probability; (iii) the autocorrelation of absolute returns with respect to lag. In contrast to previous studies, we systematically analyse all three power-laws simultaneously using the same underlying model by making observations at different time-scales. We show that the first two scaling laws are remarkably robust to the time-scale over which observations are made, irrespective of the model configuration. However, in contrast to previous studies, we show that herding may explain why long-memory is observed at all frequencies.

*Email: sphelps@essex.ac.uk

†Email: wlng@essex.ac.uk

1 Introduction

The recent financial crisis highlights the difficulties inherent in models such as CAPM which are based on rational expectations and efficient markets assumptions. Whilst it would be an exaggeration to state that belief in the efficient markets hypothesis actually caused the crisis (Ball, 2009; Brown, 2011), it is nevertheless now acknowledged that widely adopted theoretical models which assume that returns are Gaussian, as per geometric Brownian motion, are not consistent with the data from real-world financial exchanges (Lo and MacKinlay, 2001).

This had led to a resurgent interest in alternatives to models based on rational expectations models and the efficient markets hypothesis; Lo (2005) 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” (Arthur, 1994) in which it is not clear that a rational expectations solution is coherent.

Agent-based models address these issues by modeling the system in a bottom-up fashion; in a typical agent-based model we simulate the behaviour of the participants in the market — the agents — and equip them with simple adaptive behaviours. Within the context of economic and financial models, agent-based modeling can be used to simulate markets with a sufficient level of detail to capture realistic trading behaviour and the nuances of the detailed micro-structural operation of the market: for example, the operation of the underlying auction mechanism used to match orders in the exchange (Iori and Chiarella, 2002).

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. A recent area of research within the multi-agent systems community (Rayner et al., 2013; Palit et al., 2012; Cassell and Wellman, 2012) attempts to take a reverse-engineering approach

in which agent-based models of markets are built to replicate specific statistical properties that are universally observed in real-world data sets across different markets and periods of time — that is, the stylised facts of financial asset returns ([Cont, 2001](#)).

In this paper, we use an agent-based model to explore whether adaptation could be responsible for some of the *scaling laws* observed in empirical financial time-series data. 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. [Glattfelder et al. \(2011\)](#), for instance, illustrated on high-frequency foreign exchange data how 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, benefiting from the scale invariance property of the scaling law (see also [Ng, 2012](#)).

The existence of these scaling-laws presents a mystery. Consider the fact that expected absolute price changes are a power-law of the time over which they are measured. This is highly surprising since the underlying economic causes of changes to the price are very different depending on the time-horizon over which observations are made: at the highest frequency price changes are driven by the operation of the limit-order book and low-latency algorithmic trading, whereas over the longer time periods the interplay of the supply and demand of large institutional-investors is the dominant factor. Thus it is remarkable that financial time-series data should exhibit such striking self-similarity at time-scales varying from seconds to months, particularly since at higher frequencies we know that prices do not follow geometric Brownian motion.

Although there has been great deal of literature documenting the existence of these scaling laws in empirical data, to date that has been limited success in building models which explain how these scaling phenomena arise. Existing agent-based models have had

some success in showing that adaptive expectations can give rise to some of these scaling laws. For example, [LeBaron and Yamamoto \(2007\)](#) introduced an agent-based model in which power-law decay in the autocorrelation of volatility and order-signs can only be replicated when agents learn from each others' forecasts via imitation, leading to herding in expectations. When their model is analysed under a treatment in which this herding does not occur, the corresponding long-memory properties disappear. However, this is not conclusive evidence that these scaling properties arise from learning or herding behaviour; [Tóth et al. \(2011\)](#) examine long-memory in order-signs by comparing two different models, and show that long-memory in order signs could be better explained as arising from the splitting of large institutional orders into smaller orders in order to mitigate market impact — that is, it is likely caused by an exogenous influence on the market rather than arising endogenously.

Nevertheless there is not necessarily a single unified cause of all forms of scaling, and there remains the possibility that endogenous factors could account for other scaling laws. Although scaling laws relating to autocorrelation and size distributions have received a great deal of attention from the agent-based modelling community, there are relatively few models which attempt to explain other forms of scaling, such as the power-law relating absolute price changes to the time-period over which they are measured. Indeed there has been little attention paid to the fact that many scaling-laws continue to hold irrespective of the time-period over which the quantities of interest are measured. In this paper we take a systematic approach to exploring such phenomenon by examining whether the same underlying model is able to simultaneously reproduce three commonly-reported power-law relationships at different time-scales.

The outline of this paper is as follows. In the following section we give an introduction to scaling laws in finance, and the particular forms of scale invariance that we analyse in this paper. In [Section 3](#) we describe the limit-order book agent-based model that we use for our analysis. We present our findings in [Section 4](#). Finally, we conclude in [Section 5](#).

2 Scaling Laws

In Section 2.1, we first briefly explain the concept of scale invariance. Many phenomena in financial markets can be described by scaling laws. One of them studies the change in the variation of the times series (e.g. volatility of returns, absolute value of returns) in relation to the sampling frequency at which the time series is observed. This type of analysis aims to extract the scaling exponent related to the fractal dimension and to the so-called Hurst exponent. Using this popular scaling law as an example, Section 2.2 describes the estimation and a generalisation of this model in detail. Further two scaling laws on the power-law decay of the autocorrelation function and probability density function of absolute returns are then briefly presented in Section 2.3.

2.1 Scale Invariance

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

$$f(x) = ax^b \tag{1}$$

that satisfies the property of scale invariance

$$f(cx) \propto f(x) \tag{2}$$

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:

$$f(cx) = a(cx)^b = c^b ax^b = c^b f(x) \propto f(x). \tag{3}$$

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

$$\log(f(x)) = \log(a) + b \log(x). \tag{4}$$

This scale invariance of $f(x)$ is a main feature of scaling laws, allowing the researcher to model complex phenomena in a very efficient manner ([Glattfelder et al., 2011](#)).

For example, consider a standard Random Walk

$$X_t = X_{t-1} + \varepsilon_t = \sum_{i=0}^t \varepsilon_i \quad (5)$$

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 (a multiple of Δt), here denoted as $\mu_X(t)$ and $\sigma_X^2(t)$, can be straightforwardly calculated as a multiple of the innovations' moments

$$\mu_X(t) = E\left(\sum_{i=0}^t \varepsilon_i\right) = \sum_{i=0}^t E(\varepsilon_i) = t\mu_\varepsilon \quad (6)$$

$$\sigma_X^2(t) = Var\left(\sum_{i=0}^t \varepsilon_i\right) = \sum_{i=0}^t Var(\varepsilon_i) = t\sigma_\varepsilon^2. \quad (7)$$

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. (7) as

$$\sigma_X(t) = \sigma_\varepsilon \cdot t^{1/2} \quad (8)$$

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

2.2 Scaling Exponent

Let $r_{\Delta t}$ denote the log-return of a price series sampled at time interval Δt (see Figure 1 and eq. (21)). The interval Δt is a continuous variable. We follow the common approach in the econophysics literature, choose a base of 2 and consider the series of its powers in increasing order, i.e. $2^1, 2^{1.5}, 2^2, 2^{2.5}, \dots, 2^9$. If there exists a scaling law for the underlying data generating process, it will hold for any real valued base. To obtain a range of suitable

Figure 1: The top panel shows an extract of a sample price series. 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 frequencies $\Delta t = 16$ (blue) and $\Delta t = 91$ (red), respectively. The distribution of the corresponding returns is illustrated in the histograms in the bottom panels. As expected, a higher sampling interval reduces the amount of (close to) zero returns and increases the amount of extreme returns.

discrete sampling intervals, we simply round the dyadic series, i.e. $2 = 2^1, 2^{1.5} \approx 3, 4 = 2^2, 2^{2.5} \approx 6, \dots$, yielding the values $\{2, 3, 4, 6, \dots, 512\}$. These values represent discrete time steps in our agent-based model. The maximum value for Δt is set to 512 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 discrete time intervals. Figure 1 illustrate this procedure with $\Delta t = 16$ and $\Delta t = 91$ as examples. The resulting return distributions (bottom panels) do not seem to differ much in their characteristics but are 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 Δt only changes the factor of proportionality but not the functional relationship itself.

One popular scaling law models the size of the average absolute price change $E(|r_{\Delta t}|)$ as a linear function of the time interval Δt of its occurrence

$$E(|r_{\Delta t}|) \propto \Delta t^b, \quad (9)$$

where b refers to the Hurst exponent that is often estimated in long range dependence analysis. Eq. (9) is particularly interesting for applied risk management as it suggests that the expected absolute price change is proportional to the elapsed time interval raised to a power b .

Consider eqs. (1) and (9). Now define the absolute return as

$$X_{\Delta t} = |r_{\Delta t}| \quad (10)$$

and its mean as

$$E(X_{\Delta t}) = \langle X_{\Delta t} \rangle. \quad (11)$$

Note that $\langle X_{\Delta t} \rangle$ is a random variable itself and depends on the sampling interval Δt (see also eqs. (6) and (7)). If the relationship

$$\langle X_{\Delta t} \rangle = a(\Delta t)^b \quad (12)$$

holds, it would follow that

$$E(X_{c\Delta t}) = c^b E(X_{\Delta t}) \propto E(X_{\Delta t}), \quad (13)$$

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 (Ng, 2012), *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. (12)

$$\log(\langle X_{\Delta t} \rangle) = \tilde{a} + b \log(\Delta t) \quad (14)$$

with $\tilde{a} = \log(a)$ (see also eq.(4)).

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 in this log-log model measures the percentage change in the average absolute price change with respect to the percentage change of the considered time interval. For example, if the estimate for b is positive and significant, the expected price change grows with the elapsed time (diverging behaviour, similar to a Random Walk). If b is equal to zero, the variance of the underlying price process will not depend on the time frame, similar to a White Noise process. If b is negative, the expected price change decreases as time grows and would exhibit convergence to a constant value. The latter two cases, however, are less realistic and have not been reported for empirical prices processes.

Once we have calculated $\langle X_{\Delta t} \rangle$ for all considered sampling intervals (see Figure 2, right panel), we estimate the scaling law parameters according to eq. (14). The intercept and slope of the green line in the right panel in Figure 2 correspond to parameters $\tilde{a}$ and b in eq. (14).

Many empirical studies have reported that financial asset returns do not have Gaussian distributions but exhibit “stylised facts” (Cont, 2001; Lux, 2009). It is well-known that

Figure 2: The left panel compares boxplots of the *absolute* returns sampled at the different sampling intervals Δt , exhibiting their ‘self-similarity’ at different timescales. The distributions of the absolute values of the returns shown in Figure 1 are represented here in the blue and red boxplots. For each sampling frequency, consider only the means of the absolute return distributions. The parameters of the corresponding scaling law in eq. (16) are then obtained by regressing these averages against the sampling intervals Δt in the log-log dimension.

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 (7) will become more accurate because of error cancellation and volatility mean reversion, allowing for more robust estimates (McAleer and Medeiros, 2008).

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. (9) that considers higher (and possibly non-integer) moments of the absolute price change can be specified as

$$(E(|r_{\Delta t}|^{pow})) \propto \Delta t^b, \quad (15)$$

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. (15) case with *pow* = 2,

Figure 3: Consider the means of the absolute returns raised to the power $pow =$ (left panel) or $pow = 4$ (right panel). The parameters of the corresponding scaling law in eq. (16) are then obtained by regressing these averages against the sampling intervals Δt in the log-log dimension.

likewise eq. (8).

In contrast to the financial literature which mostly only consider the special cases $pow = \{1, 2\}$ (Glattfelder et al., 2011), we also consider a generalisation similar to eq. (15), also to assess the robustness of the scalability of our results. In particular, we are interested in the relationship

$$\langle X_{\Delta t}^{pow} \rangle = a_{pow}(\Delta t)^{b_{pow}}, \quad (16)$$

where pow refers to the power of the moment of interest, e.g. $\langle X_{\Delta t}^2 \rangle = E(|r_{\Delta t}|^2)$ (see eq. (8) and Figure 3, left panel). If we were interested in the kurtosis of the returns distribution, i.e. their fourth moment, to assess their “fat tails” and extreme events, we can consider eq. (16) and estimate the corresponding parameters $\tilde{a}_4$ and b_4 . As can be seen in the right panel in Figure 3, a scaling law relationship can also be found for the quartic order. This demonstrates again that the scale invariance of scaling laws (16) 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.

2.3 Other Scaling Laws

In this section, we briefly introduce a two other popular scaling laws in the econophysics literature ([Bouchaud et al., 2009](#); [Glattfelder et al., 2011](#); [Plerou et al., 2001](#)). For a more comprehensive overview, we refer the interested reader to the surveys by [Bouchaud \(2001, 2002\)](#) and [Mandelbrot \(2001a\)](#), and the references therein.

The two other scaling laws that have been widely studied in econophysics and computational finance relate to power-law *decay in autocorrelations* (e.g. long-memory in volatility) on the one hand, and power-law *size distributions* on the other (e.g. the magnitude of returns). The former is usually considered to study time dependence of a stochastic process and in particular its persistence (see also [Bouchaud et al., 2004](#))[Bouchaud et al. \(2006\)](#)[Lillo and Farmer \(2004\)](#). Let $Corr(X_{\Delta t}, X_{\Delta t-\tau}) = \langle X_{\Delta t}, X_{\Delta t-\tau} \rangle$ denote the autocorrelation function of $X_{\Delta t}$ with lag τ , then we specify

$$\langle X_{\Delta t}, X_{\Delta t-\tau} \rangle = a(\tau)^b. \quad (17)$$

The latter scaling law focuses on the tail region of the probability distribution of the return that can be modelled by power-law decay rather than an exponential decay as specified in the ubiquitous normal distribution:

$$Prob(X \geq x) = a(x)^b \quad (18)$$

The scaling exponent b in eq. (18) is often referred to as the tail index, its estimation helps to assess the frequency and magnitude of extreme events. All scaling laws presented in this section have the same functional relationship as described in eq. (12), so we estimate their scaling exponents by regression, as shown in eq. (14).

Figure 4: Consider distribution of *absolute* returns sampled at $\Delta t = 16$ (see also Figure 2, left panel). The left panel plots the tail region of the log-density against the log values of the corresponding extreme returns (see eq. (18)). The right panel shows an ACF plot of the time series of the absolute returns in logarithmic scales (see eq. (17)).

3 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 (Iori and Chiarella, 2002; LeBaron and Yamamoto, 2007; Rayner et al., 2013)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; i.e. the model incorporates adaptive expectations (Lo, 2005).

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

¹ 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 artefacts 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.

specified price and quantity. Orders are matched using a continuous double auction (CDA) ([Friedman and Rust, 1993](#)).

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 denoted 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 agents' decision problem is decomposed into two separate components: (i) calculation of their expectations of the future price of the asset (which determines their valuation thereof); and (ii) their execution strategy, which determines how they place bids or asks in the market as a function of their valuation.

The expectations component is based on an “adaptive expectations” framework ([Lo, 2005](#)) in which agents make trading decisions based on forecasts of the next period price, which are then updated through a social learning process. Agents' expectations are modelled according to a chartist, fundamentalist and noise framework ([Lux and Marchesi, 1999](#)). 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 and 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. 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, a 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 execution 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 [Gode and Sunder \(1993\)](#); every agent maintains a randomly chosen markup $\delta_i \sim U(0, \delta_{max})$ and the price of the order is set to

$$\rho_{(i,t)} = v_{(i,t)}(1 + \delta_i) \quad (19)$$

for a sell order, or

$$\rho_{(i,t)} = v_{(i,t)}(1 - \delta_i) \quad (20)$$

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:

$$r_{j,\Delta t} = \log(p_{j\Delta t}) - \log(p_{(j-1)\Delta t}) \quad (21)$$

where p_t is the market price observed at time $t (= j \times \Delta t)$, and Δt 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.

3.1 Valuations

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

$$\hat{rc}_{(i,t)} = \sum_{j=t/\Delta t - w_i}^{t/\Delta t} r_{j,\Delta t} / w_i \quad (22)$$

where w_i is the window size used by agent i . Agents using a fundamentalist forecasting rule (rf) have information about the fair value of the asset p_f and forecast accordingly:

$$\hat{rf}_{(i,t)} = \log(p_t) - \log(p_f). \quad (23)$$

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

$$\hat{rn}_{(i,t)} = \epsilon_t \quad (24)$$

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

As in [Iori and Chiarella \(2002\)](#), 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

$$\mathbf{R}_{(i,t)} = (\hat{rf}_{(i,t)}, \hat{rc}_{(i,t)}, \hat{rn}_{(i,t)}) \quad (25)$$

denote the vector of return forecasts made according to each of the fundamentalist, chartist and noise rules respectively. Let $\mathbf{S}_{(i,t)} \in \mathbb{R}^3$ 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:

$$\hat{r}_{(i,t)} = \mathbf{S}_{(i,t)} \cdot \mathbf{R}_{(i,t)}. \quad (26)$$

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

$$v_{(i,t)} = p_t e^{\hat{r}_{(i,t)} \tau_i}. \quad (27)$$

This can then be substituted into eqs. (19) and (20) in order to determine the agent's limit price.

Note that by design, the fundamental price p_f in eq. (23) is a fixed constant for all agents. Thus all innovations to prices occur endogenously as a result of agents' expectations. We do not impose any external shocks to the fundamental price, and thus any observed scaling behaviour can be attributed to the model itself and not some exogenous stochastic process.

3.2 Initial conditions

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

$$\begin{aligned}\text{SF} &\sim |N(0, \sigma_f)|, \\ \text{CF} &\sim N(0, \sigma_c), \\ \text{SN} &\sim |N(0, \sigma_n)|\end{aligned}\tag{28}$$

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.

3.3 Learning

As in [LeBaron and Yamamoto \(2007\)](#) and [Rayner et al. \(2013\)](#), we model adaptive expectations [Lo \(2005\)](#) by allowing agents to learn their forecasting strategy $\mathbf{S}$ 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:

$$\bar{\eta}_{(i,t)} = \alpha(\hat{r}_{(i,t)} - r_{(i,t)})^2 + (1 - \alpha)\bar{\eta}_{(i,t-w_\eta)}.\tag{29}$$

The fitness of agent i is then given by

$$\phi_{(i,t)} = 1/(1 + \bar{\eta}_{(i,t)}). \quad (30)$$

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:

$$p(\text{agent}_{j \neq i}) = \phi_{(j,t)} / \sum_{k \neq i} \phi_{(k,t)}, \quad (31)$$

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

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 $\mathbf{S}$ from the distributions specified in Section 3.2.

3.4 Simulation

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 3.3. For each treatment we run total of 200 independent simulations with free parameters drawn *i.i.d.* from the distributions specified in Table 1. We then analyse the scaling properties of the resulting price time-series as described in Section 2.

4 Results

In the following, we estimate the parameter sets $(\tilde{a}_{pow}, b_{pow})$ for the scaling law (16) with $pow = \{1, 1.25, 1.5, \dots, 5\}$. This procedure is then repeated for all the price series

Symbol	Definition
$p_0 \sim U(100, 1000)$	Initial price
$n \sim U(80, 150)$	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
$\tau_i \sim U(5, 10)$	Forecast time horizon
$w_\eta \sim U(100, 200)$	Sampling interval for forecast errors
$w_i \sim U(1, 100)$	Chartist window size
$\lambda_m \sim U(0.001, 0.2)$	Mutation probability per agent per time step
$\mu_f \sim U(0, 0.1)$	Drift of fundamental price process
$\sigma_f \sim U(0, 1.0)$	Volatility of fundamental price process

Table 1: Summary of model parameters and their associated distributions

produced by each of the models that we analyse. The methodology is briefly summarised in Algorithm 1. To estimate the scaling exponents for eqs. (17) and (18), we follow the similar approach as outlined in Algorithm 1, but without the consideration of the power of the moment. In the following section we describe an agent-based model of a financial market which exhibits the scale invariance under time dilation that we have described in this section.

In both agent-based models, the vast majority of the simulated time series have pronounced scaling properties (see example in Figure 2), because almost all estimated coefficients ($\widehat{\tilde{a}_{pow}^k}, \widehat{b_{pow}^k}$) (see eqs. (14) and (16)) are different from zero at the 1% significance level, and the corresponding R^2 from the regression is larger than 95%.

Figure 5 summarises our results obtained for the scaling law (16). The top left panel shows the distribution of the 200 estimated scaling exponents $\widehat{b_{pow}^k}$ for both the learning and the static model. As it can be seen in the apparent linear structure of the boxplots, this scaling law exhibit a unifractal relationship across all considered moments (pow), in both agent-based models. The distributions of the corresponding regression statistics are shown in the remaining panels. The t -values (bottom left) are all far higher than 1.96, the

Algorithm 1: Estimate scaling law $\langle X_{\Delta t}^{pow} \rangle = a_{pow}(\Delta t)^{b_{pow}}$

Data: Price Series p_t

Result: $(\tilde{a}_{pow}^k, \tilde{b}_{pow}^k)$ of all price series for both data generating processes

Consider the models (see Section 3.4):

1. learning
2. static

begin

Select model $m = \{1, 2\}$

for sample price process $k = 1$ **to** $k = 200$ **do**

for power $pow_1 = 1$ **to** $pow_{max} = 5$ **do**

for time scale $\Delta t = 2$ **to** $\Delta t = 512$ **do**

Sample p_t at $\Delta t = \{0, \Delta t, 2\Delta t, 3\Delta t, \dots\}$

Calculate $X_{\Delta t}$ (see eqs. (10) and (21))

Calculate $\langle X_{\Delta t}^{pow} \rangle = E((X_{\Delta t})^{pow})$ (see eqs. (15) and (16))

Regress $\log(\langle X_{\Delta t}^{pow} \rangle) = \tilde{a}_{pow} + b_{pow} \log(\Delta t)$ (see eqs. (14) and (16))

Save estimates $(\tilde{a}_{pow}^k, \tilde{b}_{pow}^k)$, their t -values and R^2 or regression

end

end

end

end

Compare estimated scaling exponents $\tilde{b}_{pow}^k$ and other regression measures for the two models

critical value at the 95% significance level. In all cases there is remarkably robust scaling for all moments and for all time-scales, irrespective of whether agents learn or not; for the most part the R^2 and the standard error of the regression (SER), depicted in the top right and bottom right panels respectively, remain acceptable irrespective of the sampling interval or moment we are considering.

Figure 6 summarises the results for the scaling law (18), characterising the power law decay of the probability density function in the tail region. The findings are very similar to above. Across all time scales, the scaling exponent is in the learning model higher than in the static model, implying a slower decay of the density function, and hence a higher probability of extreme returns – compared to the static model. However, as the sampling interval increases, the scaling exponent decreases, i.e. the tails are becoming less ‘fat’. In the majority of cases, the absolute t -values and R^2 s are very high, both implying a good fit of the model in general. Except for the very short time intervals, a comparison of the

Figure 5: Consider the 200 regressions of the scaling law model $\langle X_{\Delta t}^{pow} \rangle = a_{pow}(\Delta t)^{b_{pow}}$ (see also Figures 2 and 3) for the learning and the static model. These boxplots illustrate the distributions of the scaling exponents (top left) and their corresponding t -values (bottom left), and the R^2 s (top right) and standard errors (bottom right) of the regressions.

SERs show that the residual errors of the fit are mostly similar.

We also follow the approach in Plerou et al. (2001) and re-estimate scaling law (18) for the left (r^{neg}) and right (r^{pos}) tail of the return distribution in order to study its (a)symmetry. This is to ensure that the observed scaling behaviour in the fat tail distribution of $X_{\Delta t}$ does not arise from one particular side of the distribution or $r_{\Delta t}$. The results are illustrated in Figure 7. As discernible, the scaling pattern is linear and unifractal for both tails, and very similar to the scaling observed for $x_{\Delta t}$. Similarly to Plerou

Figure 6: Consider the 200 regressions of the scaling law model $Prob(X_{\Delta t} \geq r) = a(r)^b$ (see also Figure 4, left panel) for the learning and the static model. These boxplots illustrate the distributions of the scaling exponents (top left) and their corresponding t -values (bottom left), and the R^2 s (top right) and standard errors (bottom right) of the regressions.

et al. (2001), we also applied the Hill estimator to both tails. The boxplots in the bottom panels generally reconfirm the findings obtained so far. The Hill estimates are very similar for both tails, implying a symmetric distribution. The values for the learning models are generally lower (than those of the static ones), indicating a fatter tail.

In the two scaling models we discussed so far, the fit of the power-law is resilient to changes in the free parameters of our model. Although learning does not dictate the fit of these scaling-laws (i.e. whether or not they are present), it does increase the exponent;

Figure 7: These subplots analyse whether the positive (r^{pos}) and negative returns (r^{neg}) have different tail indices. Consider the 200 regressions of the scaling law models $Prob(r_{\Delta t}^{pos} \geq r) = a(r)^b$ (left panels) and $Prob(|r_{\Delta t}^{neg}| \geq r) = a(r)^b$ (right panels) for the learning and the static model. The boxplots illustrate the distributions of the scaling exponents (top panels). As can be seen, the return distributions are relatively symmetric, also indicated by the corresponding Hill-estimates displayed in the bottom panels.

learning leads to fatter tails and greater price dispersion. These results are consistent with an analysis of empirical data which finds that larger scaling exponents $b > 0.5$ are associated with emerging markets, whereas lower scaling exponents are associated with indices representative of more mature markets, such as the S&P 500 (Di Matteo et al., 2005). The authors conjecture that this might be explained by the heterogeneity of agents' expectations; in more mature markets we might expect less herding, and a corresponding

Figure 8: Consider the 200 regressions of the scaling law model $\langle X_{\Delta t}, X_{\Delta t-\tau} \rangle = a(\tau)^b$ (see also Figure 4, right panel) for the learning and the static model. These boxplots illustrate the distributions of the scaling exponents (top left) and their corresponding t -values (bottom left), and the R^2 s (top right) and standard errors (bottom right) of the regressions.

lower exponent. Our analysis corroborates this picture.

As with other agent-based models of financial markets, we find that as agents herd they induce short-lived trends in the price — i.e. speculative bubbles and crashes. The striking feature of our analysis in contrast to previous studies is that this effect of herding is unifractal (see also Calvet and Fisher, 2002; Lux, 2009). On the other hand, the power-law in the autocorrelation of absolute returns (eq. (17)) is more subtle. Both the fit and the exponent are sensitive to the configuration of the model. In contrast to

LeBaron and Yamamoto (2007), we find that for higher-frequencies we observe power-law autocorrelation of absolute returns both with and without learning. However, the scaling exponent for the learning models are generally higher than in the static model, implying a slower decay of the autocorrelation function, and hence a higher persistence as in the static model. As we increase the sampling interval, we see that long-memory is present at all time-scales for the learning case, whereas it gradually disappears when agents do not learn (compare the *static* and *learning* boxplots in the top right panel of Figure 8). Furthermore, although the R^2 is similar for both models at the higher sampling frequencies, the SERs are always higher for the static model (bottom right).² In our model, it is not simply that imitation results in long-memory, but that herding gives rise to a temporal-symmetry of this property.

Other research has highlighted the importance of the *dynamics* of learning (Rayner et al., 2013) in reproducing particular statistical features observed in empirical data, and (Hommes and Sorger, 1998; Sato et al., 2002; Skyrms, 1992) 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 (Hommes et al., 2008). We conjecture that the unifractal long-memory behaviour exhibited in our current study is the result of similar chaotic attractors in the co-evolutionary dynamics resulting from learning; oscillations in the dynamics of learning occur at many different frequencies, yielding long-range dependence at many different time-scales. An analysis of the underlying dynamics of learning in order to determine its fractal nature promises to be an interesting area of future study.

²Note that we have not considered sampling intervals higher than $\Delta t = 32$ as the autocorrelation function became negative in some of the longer lags.

5 Conclusion

In this paper, we modelled the financial market using an agent-based model in which agents' expectations are driven by *heuristic* forecasting rules, and in which agents are able to imitate each other via a social-learning process, i.e. a form of herding behaviour. We have shown that within this framework, we can reproduce *unifractal* scaling with respect to three well-known power-laws relating: (i) moments of the absolute price-change to the time-scale over which they are measured; (ii) the magnitude of returns with respect to their probability; (iii) the autocorrelation of absolute returns with respect to the lag.

In contrast to previous studies, we systematically analysed all three power-laws simultaneously using the same underlying model by making observations at different time-scales. We showed that the first two scaling laws are remarkably robust to the time-scale over which observations are made, irrespective of the model configuration. However, in contrast to previous studies, we showed that herding may explain why long-memory is observed over all time-scales.

In this study, we only considered the scaling behaviour of absolute prices changes and returns. 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, as well as the underlying learning dynamics themselves, to study their fractal nature and degree of self-similarity.

Acknowledgements

The authors are grateful to Khurom Kiyani and Imon Palit for their valuable and helpful discussions. The authors also would like to thank Daniel O'Leary (editor) and two anonymous referees for their constructive comments and suggestions that led to an improvement of the paper.

References

- Arthur, W. B. (1994). Inductive Reasoning and Bounded Rationality. *The American Economic Review*, 84(2):406–411.
- Ball, R. (2009). The Global Financial Crisis and the Efficient Market Hypothesis: What Have We Learned? *Journal of Applied Corporate Finance*, 21(4):8–16.
- Bouchaud, J.-P. (2001). Power laws in economics and finance: some ideas from physics. *Quantitative Finance*, 1(1):105–112.
- Bouchaud, J.-P. (2002). An introduction to statistical finance. *Physica A: Statistical Mechanics and its Applications*, 313(1-2):238–251.
- Bouchaud, J.-P., Farmer, J., and Lillo, F. (2009). How markets slowly digest changes in supply and demand. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editor, *Handbook of Financial Markets: Dynamics and Evolution*, chapter 2, pages 57–160. North-Holland, San Diego.
- Bouchaud, J.-P., Gefen, Y., Potters, M., and Wyart, M. (2004). Fluctuations and response in financial markets: the subtle nature of ‘random’ price changes. *Quantitative Finance*, 4(2):176–190.
- Bouchaud, J.-P., Kockelkoren, J., and Potters, M. (2006). Random walks, liquidity molasses and critical response in financial markets. *Quantitative Finance*, 6(2):115–123.
- Brown, S. J. (2011). The Efficient Markets Hypothesis: The Demise of the Demon of Chance? *Accounting and Finance*, 51(1):79–95.
- Calvet, L. and Fisher, A. (2002). Multifractality in Asset Returns: Theory and Evidence. *Review of Economics and Statistics*, 84(3):381–406.
- Cassell, B.-A. and Wellman, M. P. (2012). Asset pricing under ambiguous information: an empirical game-theoretic analysis. *Computational & Mathematical Organization Theory*, 18(4):445–462.

- Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236.
- Di Matteo, T., Aste, T., and Dacorogna, M. M. (2005). Long-term memories of developed and emerging markets: Using the scaling analysis to characterize their stage of development. *Journal of Banking & Finance*, 29(4):827–851.
- Friedman, D. and Rust, J., editors (1993). *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.
- Glattfelder, J. B., Dupuis, A., and Olsen, R. B. (2011). Patterns in high-frequency FX data: discovery of 12 empirical scaling laws. *Quantitative Finance*, 11(4):599–614.
- Gode, D. K. and Sunder, S. (1993). Allocative efficiency of markets with zero-intelligence traders: Market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.
- Hommes, C., Sonnemans, J., Tuinstra, J., and van de Velden, H. (2008). Expectations and bubbles in asset pricing experiments. *Journal of Economic Behavior & Organization*, 67(1):116–133.
- Hommes, C. and Sorger, G. (1998). Consistent Expectations Equilibria. *Macroeconomic Dynamics*, 2(3):287–321.
- Iori, G. and Chiarella, C. (2002). A Simulation Analysis of the Microstructure of Double Auction Markets. *Quantitative Finance*, 2(5):346–353.
- LeBaron, B. and Yamamoto, R. (2007). Long-memory in an order-driven market. *Physica A: Statistical Mechanics and its Applications*, 383(1):85–89.
- Lillo, F. and Farmer, J. D. (2004). The Long Memory of the Efficient Market. *Studies in Nonlinear Dynamics & Econometrics*, 8(3):Article 1.

- Lo, A. W. (2005). Reconciling Efficient Markets with Behavioural Finance: The Adaptive Markets Hypothesis. *Journal of Investment Consulting*, 7(2):21–44.
- Lo, A. W. and MacKinlay, A. C. (2001). *A Non-Random Walk Down Wall Street*. Princeton University Press, Princeton, NJ.
- Lux, T. (2009). Stochastic behavioral asset pricing models and the stylized facts. In Thorsten Hens and Klaus Reiner Schenk-Hoppé, editor, *Handbook of Financial Markets: Dynamics and Evolution*, chapter 3, pages 161–215. North-Holland, San Diego.
- Lux, T. and Marchesi, M. (1999). Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(February):498–500.
- Mandelbrot, B. (2001a). Scaling in financial prices: I. Tails and dependence. *Quantitative Finance*, 1(1):113–123.
- Mandelbrot, B. (2001b). Scaling in financial prices: II. Multifractals and the star equation. *Quantitative Finance*, 1(1):124–130.
- McAleer, M. and Medeiros, M. (2008). Realized Volatility: A Review. *Econometric Reviews*, 27(1):10–45.
- Ng, W. L. (2012). Improving Long-Term Financial Risk Forecasting Using High-Frequency Data and Scaling Laws. In Alexandrova-Kabadjova, B., Martinez-Jaramillo, S., Garcia-Almanza, A. L., and Tsang, E., editors, *Simulation in Computational Finance and Economics: Tools and Emerging Applications*, pages 255–278. IGI Global.
- Palit, I., Phelps, S., and Ng, W. L. (2012). 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. International Foundation for Autonomous Agents and Multiagent Systems.

- Plerou, V., Gopikrishnan, P., Gabaix, X., Amaral, L., and Stanley, H. (2001). Price fluctuations, market activity and trading volume. *Quantitative Finance*, 1(2):262–269.
- Rayner, N., Phelps, S., and Constantinou, N. (2013). Testing adaptive expectations models of a double auction market against empirical facts. In *Lecture Notes on Business Information Processing: Agent-Mediated Electronic Commerce Designing Trading Strategies and Mechanisms for Electronic Markets*, pages 44–56. Springer, Barcelona.
- Sato, Y., Akiyama, E., and Farmer, J. D. (2002). 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.
- Skyrms, B. (1992). Chaos in Game Dynamics. *Journal of Logic, Language, and Information*, 1:111–130.
- Tóth, B., Palit, I., Lillo, F., and Farmer, J. D. (2011). Why is order flow so persistent? Technical Report arXiv:1108.1632v1.