

MULTIFRACTAL ANALYSIS OF AGENT-BASED FINANCIAL MARKETS

James R. Thompson
James R. Wilson

Edward P. Fitts Department of Industrial and Systems Engineering
North Carolina State University
Raleigh, NC 27607, USA

ABSTRACT

To analyze financial time series exhibiting volatility clustering, long-range dependence, or heavy-tailed marginals, we exploit multifractal analysis and agent-based simulation. We develop a robust, automated software tool for extracting the multifractal spectrum of a time series based on multifractal detrended fluctuation analysis (MF-DFA). The software is tested on simulated data with closed-form monofractal and multifractal spectra to ensure the quality of our implementation. We perform an in-depth analysis of General Electric's stock price using traditional time series techniques, and contrast the results with those obtained using MF-DFA. We also present a zero-intelligence agent-based financial market model and analyze its output using MF-DFA. We study the changes in the macrolevel time series output as analyzed by MF-DFA when altering one of the microlevel agent behaviors. Finally we explore the potential for validating agent-based models against empirical time series using MF-DFA.

1 INTRODUCTION

The traditional approach for modeling systems with a large number of stochastic variables, such as financial markets, has been discrete-event simulation (DES). But as the availability of computing power and data storage capacity have increased, the concept of agent-based modeling has become more appealing (North and Macal 2007). DES models focus on the processes of a system — that is, the sequence of events “seen” by homogeneous agents (or entities) as those agents move through the system and are impacted by the events. By contrast, agent-based models focus on agent behavior, allowing the agents a degree of influence over which events they experience. Moreover, the agents' reactions to system events are heterogeneous, and the impact of microlevel behaviors on macrolevel system output can be analyzed.

Economics, in particular, has adopted agent-based modeling to test theories of behavioral economics, the formation of prices under uncertainty, and the impact of market structure on market participants and their actions (Tesfatsion 2006). The analysis of economic systems relies heavily on time series data; similarly, the output of any agent-based model of an economic system will also require time series analysis. Many types of financial time series, most notably market returns, have been found to exhibit long-range memory as well as dramatic day-to-day swings that cannot be adequately represented by light-tailed distributions such as the normal distribution. In particular, this means that for such time series, the usual variance parameter (i.e., the sum of covariances at all time lags) is not defined because the covariance function does not converge to zero fast enough as the time lag increases. Moreover in such time series, often the tails of the marginal density converge to zero so slowly that higher-order marginal moments such as skewness and kurtosis fail to exist. Conventional autoregressive methods for analyzing time series data often fail to accommodate these properties, and modern approaches such as autoregressive conditional heteroscedastic (ARCH) models also fall short of capturing the true erratic nature of empirical data (Engle 1982). However, Farmer, Patelli, and Zovko (2005) used a “zero-intelligence” agent-based financial market simulation to

create output that exhibited volatility clustering and heavy-tailed returns simply due to the rules for trading in a double auction order book.

As early as 1900, researchers began modeling financial time series with highly irregular structures like Brownian motion. Brownian motion is known to be a self-similar process whose sample paths are continuous, nondifferentiable, and have Hölder exponent $\alpha = 0.5$ at every time $t \geq 0$. Mandelbrot's pioneering work in geometry later showed that irregular structures of this nature are so rough that they possess a Hausdorff dimension that strictly exceeds their topological dimension (Mandelbrot 1977). He called these objects "fractals" and went on to generalize the concept to more complex shapes known as multifractals. In many application domains, a multifractal time series exhibits self-similar behavior; but over each time interval $[t, t+s]$, the magnitude of the associated fluctuation scales with a local Hölder exponent $\alpha(t)$ that may vary as a function of the time t . Over a finite time horizon $[0, T]$, the renormalized density of these Hölder exponents $\{\alpha(t) : t \in [0, T]\}$ stochastically quantifies the multifractal and is known as the multifractal spectrum. Building on the idea that the multifractal spectrum can stochastically quantify a time series, Kantelhardt et al. (2002) devised an algorithm—known as multifractal detrended fluctuation analysis (MF-DFA)—to extract the multifractal properties of empirical data.

The rest of this paper is organized as follows. In Section 2 we introduce the concept of fractals and multifractals in the context of time series analysis. We discuss the mathematical background of the multifractal spectrum and present multifractal detrended fluctuation analysis as a method for its extraction from empirical data. We will also describe the software we developed to implement this analysis and give results of its performance on simulated and empirical data. In Section 3 we discuss the zero-intelligence agent-based simulation of financial markets proposed by Farmer, Patelli, and Zovko (2005). We describe our experiment of altering the agents' method for placing bid- and ask-orders, and present the different time series created by the model. In Section 4 we analyze the output of the simulation using MF-DFA and we close in Section 5 with a brief discussion of our findings.

2 MULTIFRACTAL DETRENDED FLUCTUATION ANALYSIS

In this article we restrict attention to self-similar fractals. For a self-similar fractal object, each small part of the object resembles the whole object in some essential respect (Addison 1997). For example, an increment of Brownian motion over a fixed time period $[t, t+s]$ is a suitably rescaled probabilistic replica of Brownian motion over the much smaller time period $[t, t+\eta s]$ when $0 < \eta \ll 1$; and when $\eta \gg 1$, a similar relationship holds. Formally, the increments of a random function $\{X(t) : -\infty < t < \infty\}$ are said to be self-similar with parameter $H \geq 0$ if for any $\eta > 0$ and t , $X(t+s) - X(t) \stackrel{d}{=} \eta^{-H}[X(t+\eta s) - X(t)]$, where the $\stackrel{d}{=}$ denotes equality in distribution. The parameter H is often referred to as the Hurst exponent. Examination of self-similar data sets in the context of time series analysis shows that the Hurst exponent characterizes the asymptotic behavior of the autocorrelation function of the time series. Values of H in $(0.5, 1.0)$ lead to positive autocorrelations that decay too slowly for the sum of autocorrelations over all lags to be finite. That is, processes with $0.5 < H < 1$ are long-memory processes, whereas processes with $0 \leq H \leq 0.5$ are said to exhibit short-range dependence (Mandelbrot and van Ness 1968). Time series with $0 < H < 0.5$ exhibit antipersistence in their paths, as positive changes are more likely followed by negative changes resulting in a very rough structure. When $0.5 < H < 1$, the series exhibits persistence, as successive increments in the series are more likely to be in the same direction, creating smoother paths. The Hurst exponent also quantifies the inherent roughness exhibited by these monofractals. It can be shown that the Hausdorff dimension D of the graph of a monofractal time series is related to the Hurst exponent of the underlying process by the relation $H = 2 - D$ (Graneroa et al. 2008). For example, Brownian motion has $H = \alpha = 0.5$; and every sample path of Brownian motion has Hausdorff dimension $D = 1.5$.

A multifractal time series is defined in a similar way to monofractals, except the exponent H is no longer a single scalar value, but rather a function of Hölder exponents that are changing within the time series. The Hölder exponents are related to a generalized Hurst exponent $h(q)$ that varies with different values of the

moment order q through the multifractal formalism. The difference between these two descriptive power laws is subtle, but in general the Hurst exponent is a global quantity, while the Hölder exponent is a local quantity (Coffey and Kalmykov 2006). Formally, for some stochastic process $\theta(t)$ with support $[0, T]$, we define the Hölder exponent $\alpha(t)$ at each time $t \in [0, T]$ by the asymptotic relationship $\theta(t+s) - \theta(t) = O_P(|s|^{\alpha(t)})$ for all t as $s \rightarrow 0$, where $\alpha(t) = \sup \{ \beta \geq 0 : \theta(t+s) - \theta(t) = O_P(|s|^\beta) \text{ as } s \rightarrow 0 \}$ and in general for continuous-time stochastic processes $\{A(s) : s \in \mathbb{R}\}$ and $\{B(s) : s \in \mathbb{R}\}$, the “big-oh-sub- P ” notation $A(s) = O_P[B(s)]$ as $s \rightarrow 0$ means that given an arbitrarily small $\zeta > 0$, there is a constant $M = M(\zeta)$ and a positive number $\varepsilon = \varepsilon(\zeta)$ such that $\Pr\{|A(s)| \leq M|B(s)|\} \geq 1 - \zeta$ for $|s| < \varepsilon$.

One way to think of a multifractal time series is as the union of an infinite number of monofractal time series, each characterized by a single self-similarity exponent α . However, these monofractals are interwoven throughout the support such that the set of points supporting any one monofractal constitutes a fractal set.

Definition 1 A stochastic process $\{X(t) : t \in [0, T]\}$ is *multifractal* if it has stationary increments and satisfies

$$E[|X(t)|^q] = c(q)t^{\tau(q)+1} \text{ for all } t \in [0, T] \text{ and } q \in \mathcal{Q}, \quad (1)$$

where $T > 0$ and $[0, 1] \subseteq \mathcal{Q}$.

The function $\tau(q)$ is called the *scaling function* of the multifractal process (Calvet and Fisher 2002). The scaling function is concave (see Proposition 1 of Calvet and Fisher (2002)); moreover the multifractal spectrum $f(\alpha)$ of a multifractal stochastic process $\{X(t) : t \in [0, T]\}$ with scaling function $\tau(q)$ for $q \in \mathbb{R}^1$ is given by

$$f(\alpha) = \min_q \{q\alpha - \tau(q)\}, \quad (2)$$

which is the Legendre transform of $\tau(q)$.

If for a fixed $q \in \mathbb{R}^1$ we define the set of points $\mathcal{T}_q = \{t : \alpha(t) = h(q) + q \frac{dh(q)}{dq}\}$ and if we let $\alpha = \alpha_q$ denote the common value of the Hölder exponent $\alpha(t)$ for all $t \in \mathcal{T}_q$, then the multifractal spectrum $f(\alpha)$ evaluated at $\alpha = \alpha_q$ is defined to be the Hausdorff dimension of the set \mathcal{T}_q . The function $f(\alpha)$ stochastically quantifies a multifractal. In general, the multifractal spectrum gives us two important quantifications of a time series: (i) the set of Hölder exponents, $\alpha(t)$, that define how the series scales as we measure on finer intervals; and (ii) the fractal dimension, $f(\alpha)$, of the set of points in time that scale with a particular exponent α . Although $f(\alpha)$ is not a proper probability density, in this renormalization it represents the relative concentration of points in time that exhibit a scaling exponent α . The multifractal spectrum is convex and peaks at the expected value of $\alpha(t)$. If a function is monofractal, like Brownian motion, it has a single exponent defining how it scales for all increments of time and thus its multifractal spectrum is a single point at $\alpha = H$ with $f(\alpha) = 1$. These various interpretations are collectively referred to as the multifractal formalism and lead to our intuition of the multifractal spectrum. They also lead to methods for extracting the spectrum from empirical data.

2.1 The MF-DFA Algorithm

Among the most effective methods for extracting $f(\alpha)$ from empirical data, multifractal detrended fluctuation analysis (MF-DFA) is the easiest to implement and the most robust (Kantelhardt et al. 2002). The objective of MF-DFA is to estimate the generalized Hurst exponent $h(q)$ for a range of negative and nonnegative values of q . When q is positive with a large magnitude (say between 5 and 10), the large fluctuations in the data will be accentuated and the small fluctuations will be diminished. Conversely, when q is negative with a large magnitude, the small fluctuations will be accentuated and the large fluctuations diminished. Thus, long memory in different size fluctuations (e.g. volatility clustering) can be detected.

The basic MF-DFA algorithm involves partitioning a cumulative sum of increments into segments of length s , fitting a polynomial to each segment, and then computing the average root mean square error of all the segments. This process is repeated for many values s , and the resulting averages are plotted against

s on a doubly-logarithmic scale. This is analogous to analyzing how the separation of points scales with time to derive the Hurst exponent. The slope of the doubly-logarithmic plot constitutes an estimate of the generalized Hurst exponent $h(q)$. Iteratively applying this process to a range of q th-root-mean-square error terms involving the residuals within each segment and plotting the results on a doubly-logarithmic scale yields a function of power laws rather than a single value. From the slope $h(q)$ of the fitted linear function we can calculate $\tau(q)$ via the relationship $\tau(q) = qh(q) - 1$ for $q \in \mathcal{Q}$, and then, applying the Legendre transform, extract α and $f(\alpha)$. For a detailed algorithmic statement of the MF-DFA procedure and our method for implementing that procedure, see §§2.4–2.5 of Thompson (2013).

2.2 Multifractal Analysis

We programmed the MF-DFA algorithm in a Java application that allows the user to manipulate the parameters s and q . We first tested the application on known monofractals and multifractals with closed form multifractal spectra to determine the accuracy of our implementation. We then conducted multifractal analysis on a high-frequency trade and quote (TAQ) data set of General Electric's (GE) stock price as traded on the New York Stock Exchange.

We chose GE as the asset for analysis arbitrarily. GE is a diverse conglomerate with manufacturing divisions in energy, technology infrastructure, and household and industrial products. They also engage in capital finance. GE is publicly traded on the New York Stock Exchange and is a component of the Dow Jones Industrial average and the S&P 500. Our data set covers the period 2000 to the end of 2003 and is shown in Figure 1(a). GE issued a 3-for-1 split on May 8, 2000. This is the only adjustment we made to the original data set. The date of the split is shown by the vertical line in the figure.

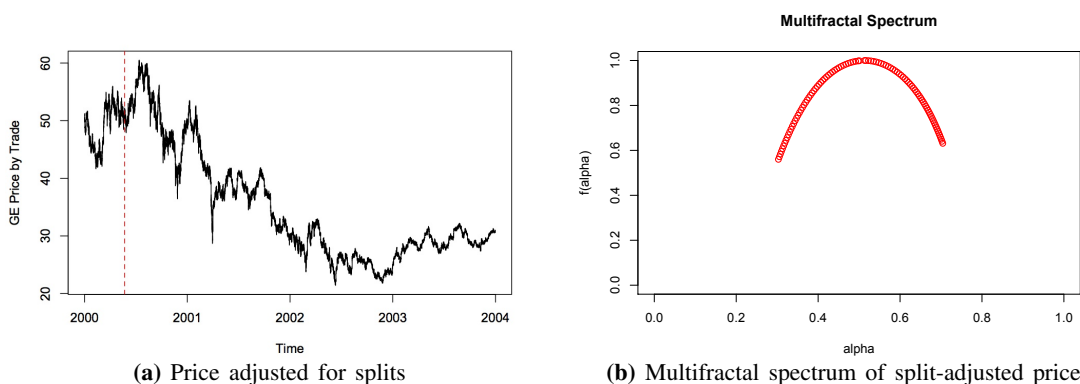


Figure 1: The fine-scale TAQ price for GE from the period 2000 to 2003.

The data set consists of 4,273,056 trades. The volume of trading steadily increased from roughly 3,200 trades per day in 2000 to over 5,300 trades per day in 2003. Results of traditional exploratory time series analysis suggested an AR(2) model would adequately model the GE series. However, the residuals of the AR(2) had extremely heavy tails, resulting in rejection of the null hypothesis of the standard Box-Ljung test. Similarly poor results were obtained from the GARCH(1,1) fit to the data set. For complete details on the AR(2) and GARCH(1,1) fits to the GE data set, see §3.3.1 of Thompson (2013)

Next we performed multifractal analysis on the GE data using MF-DFA. The multifractal spectrum extracted from the data (Figure 1(b)) gives some indication as to why the previous standard models failed to describe the statistical nature of the data set. The multifractal spectrum peaks around $\alpha = 0.48$, which indicates that the predominant fluctuations in the data are uncorrelated. But the breadth of the spectrum indicates that small fluctuations exhibit persistent (i.e. long-range dependent) behavior while large fluctuations exhibit antipersistent behavior. Both the ARMA and GARCH processes fail to capture this complex scaling behavior (Calvet and Fisher 2002).

3 THE ZERO-INTELLIGENCE MODEL

In this article, we analyze the construction of multifractal financial time series through an agent-based model. We build an agent-based model based on Farmer, Patelli, and Zovko (2005) in which the agents possess “zero intelligence,” and the double auction order book is the driving factor in the market. We adjust the probability distribution of the size of price offsets used by the agents in the model and analyze the impact on the price path and its multifractal properties.

3.1 Model Structure

The basic premise for the zero-intelligence models is that market structure, rather than agent behavior, is the dominant factor in price movements. In this case, the market structure is represented by a continuous double auction order book. That is, traders can place either buy or sell orders, and they can place these orders at any time.

Agents in the model place either limit orders or market orders. A limit order is an order placed with a price that does not currently meet the best price being offered. In the case of a buy limit order, this means the current best sell price is higher than the limit order’s bid price. In the case of a sell limit order, the opposite is true. Since the limit order terms cannot be immediately met or “crossed,” the limit order is placed in a best-price-first queue. Ties in the queue are broken on a first-come-first-serve basis. At time t , the lowest selling price $a(t)$ is the best price in the ask-order queue and the highest buying price $b(t)$ is the best price in the bid-order queue. These values are known as the boundary conditions for new limit orders. At time t , the bid-ask spread $s(t)$ is defined as the difference between the best ask price and the best bid price, $s(t) = a(t) - b(t)$. When new orders arrive, any orders that are crossed are immediately executed, and the bid-order and ask-order queues are reorganized as necessary to ensure the best buying and asking prices are first in their respective queues. Limit orders are placed by “patient” agents who are theoretically willing to wait for the market to meet their price thresholds.

Impatient agents, by contrast, are not willing to wait and therefore place market orders instead of limit orders. Market orders cause an immediate transaction. That is, buy market orders execute at the current best asking price and sell market orders execute at the current best bid price. Market orders thus create an immediate ripple in the price process as they clear the current best price, thereby changing the boundary condition for new limit orders. Patient agents place only limit orders according to a Poisson process. Impatient agents place only market orders according to another independent Poisson process. Orders of both types expire according to a third point process whose event times are derived from a Poisson process that is embedded in the order-generating process. The order expiration interevent time is set when the order is created by the patient agent, thereby making it dependent on that agent’s limit-order Poisson process. In our simplified double auction order book, all orders consist of one share of a single asset. The limit order price is constructed by offsetting the current best price (ask or bid) by a random variable generated from a selected distribution. Offsetting the current best price ensures that $a(t)$ and $b(t)$ create boundary conditions for new limit orders. Figure 2 illustrates how the order book queues limit orders as they arrive.




Figure 2: Double auction order book bid and ask queues.

We arbitrarily choose to have 50 patient agents and 50 impatient agents. Patient agents have fields for storing the next buy time and the next sell time. Each agent is instantiated and scheduled for time $t = 1$ when the simulation is started. Then each agent generates an exponential interarrival time for its next buy and sell orders, and schedules itself to place these orders accordingly. Each time thereafter when

an order is placed, a future order of the same type (buy or sell) is scheduled and the order expiration time is scheduled. Setting the price for new limit orders amounts to generating a random value Δ_p with which to offset the appropriate boundary condition. That is, for sell orders the limit order price is set to $p = b(t) + |\Delta_p|$ and for buy orders the limit order price is set to $p = a(t) - |\Delta_p|$, where the distribution of Δ_p and its impact on the multifractal properties of the price path are the subject of this study. To ensure a level of reality in the price process, limit orders with a negative price are not allowed to enter the best-price queue. The choice of distribution is predicated on stressing the discipline enforced on the market by the double auction order book. Since limit orders are arranged in the bid-order and ask-order queues according to the best-price-first queuing discipline, the market structure is such that large ask orders and small bid orders are rarely satisfied by incoming market orders. The prices placed well outside the bid-ask spread will ultimately expire before they are ever cleared. In contrast, those closer to the spread will be cleared more readily as market orders cause shifts in the current best prices. Using a thin-tailed distribution for Δ_p means that all limit orders have a price close to the spread (say, within six standard deviations); and thus the market discipline is difficult to distinguish from the general trend of the price offsets.

3.2 Experiments

The claim by Farmer, Patelli, and Zovko (2005) is that zero-intelligence models, properly calibrated to empirically observed rates of order placement, produce a good approximation to the price diffusion rate observed in actual price paths. We saw in the previous section that many financial time series exhibit a spectrum of Hölder exponents that give rise to more erratic price paths and the possibility of extreme events. Thus, we propose that MF-DFA may be more appropriate than price diffusion alone for validating an agent-based financial market against empirical data. As noted earlier, a heavy-tailed distribution for Δ_p will accentuate the impact of the double auction order book, while a thin-tailed distribution will mask the market discipline. It is thus of interest to see how the presence of thin or heavy tails in Δ_p will change the multifractal spectrum of the price path.

To investigate the impact of Δ_p on the zero-intelligence model, we chose the Pareto distribution and the Normal distribution to drive the offset random variable. The cumulative distribution function of the Pareto is given by

$$F(x) = \begin{cases} 1 - (x_{\min}/x)^c, & \text{for } x \geq x_{\min}, \\ 0, & \text{if } x < x_{\min}, \end{cases} \quad (3)$$

where $x_{\min} > 0$ represents the current best price (ask or bid). It can be shown that if $0 < c \leq 1$, then the Pareto distribution fails to have a finite mean; and for $c \in [1, 2]$ it fails to have finite variance. Bearing this in mind, we chose $c = 1.1$ for our simulation. This means the offset Δ_p has a finite mean, but an infinite variance—i.e., very heavy tails. In contrast we chose the Normal(0, 1) distribution for Δ_p to represent the thin-tailed run of the simulation. Since the Gaussian is two-sided, we took the absolute value of the random variate generated to serve as the offset Δ_p . Since we deliberately chose two extreme distributions for the parameter Δ_p , the impact on the simulated price path is profound, as shown in Figure 3 below. In the next section we apply the standard time series analysis techniques to these data sets, as well as multifractal analysis.

4 RESULTS

4.1 The Thin-Tailed Simulation

Acknowledging that these are deliberately extreme results, we seek to assess the ability of agent-based financial markets to reproduce various stylized facts observed in real financial data. Starting with the thin-tailed experiment, we note from Figure 3(b) that the price path could feasibly represent a real asset. There are 543,883 data points ranging from a low of \$0.50 to a high of \$18.79 per share. There are long periods of increase and decline, but on a finer scale there are many small fluctuations. Inspecting the first

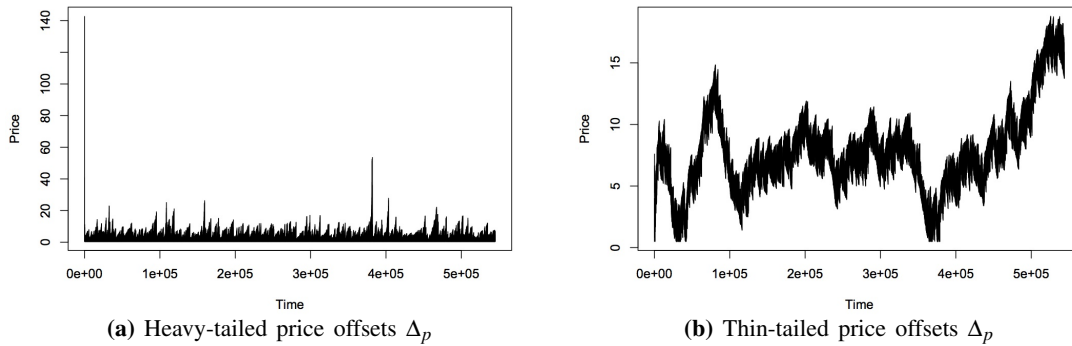


Figure 3: Zero-intelligence model with heavy- and thin-tailed price offsets.

difference of the series suggests that white noise may be governing the trade-to-trade fluctuations (see Figure 4), as the differenced series appears to be weakly stationary.




Figure 4: The first difference of the thin-tailed simulation.

Both the autocorrelation and partial autocorrelation functions of the original series decay very slowly with lag l . Given the first difference appears to be weakly stationary, we analyzed its autocorrelation structure, shown in Figure 5. The sharp decline in the ACF coupled with a slow decline in the PACF indicates that a moving average may best describe the transformed data set.




Figure 5: The ACF and PACF of the first difference of the thin-tailed simulation.

Since the ACF cuts off sharply at lag $l = 2$, we fit a moving average MA(2) model to the first difference of the price path. However, the resulting model was not statistically a good fit for the data. The Box-Ljung test indicates there is not enough evidence to conclude that the residuals are i.i.d. random variables. However, the histogram of the residuals sheds some light on why this might be the case (see Figure 6(a)). The structure of the double auction order book is such that the price changes will rarely be exactly zero. From

one time step to the next, traders are shifting the bid-ask boundary conditions causing most changes to be around zero, but never exactly zero. This creates the bimodal histogram observed in the residuals.

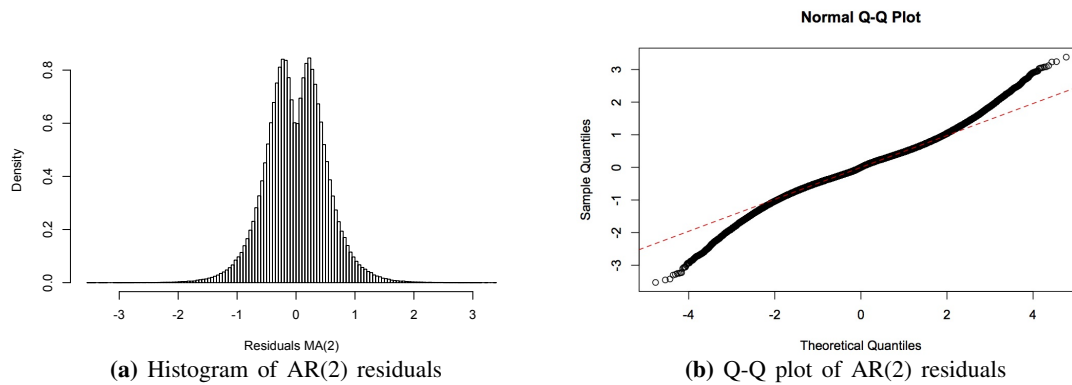


Figure 6: Residuals from the MA(2) model of the thin-tailed simulation.

Finally, we observe the QQ-plot of the residuals in Figure 6(b). The plot deviates enough from the predicted line (shown in red) to conclude the residuals are not normal random variables. However, the tails do not seem to be the reason for this deviation. Given that our price distribution for Δ_p is known to be Gaussian, we can conclude that market discipline, along with the other parameters such as order rates, combine to create an erratic price path that is poorly described by autoregressive–moving average techniques.

4.2 The Heavy-Tailed Simulation

We proceed in a similar manner for the heavy-tailed simulation. As predicted, we observe wild swings in the price as it ranges from less than a dollar to over \$140. The first difference also reveals clusters of volatility have been induced by the heavy-tailed distribution coupled with the market discipline of the order book (see Figure 7). However, the clusters appear to follow a very antipersistent pattern, as large positive changes are quickly followed by large negative changes. This clustering of volatility strongly suggests a GARCH model would be appropriate for this series.




Figure 7: The first difference of the heavy-tailed simulation.

Next we analyze the autocorrelation structure to determine if our initial thought of fitting a GARCH model is reinforced by long-term correlation. Figure 8 shows that although the correlation is small in the ACF, it continues to exceed the minimum level for significance well beyond lag $l = 50$. Similarly, the PACF fails to decay below the significance level at this lag, indicating that autoregressive moving averages

would constitute a poor fit for this data set. Pairing that with our observation of volatility clustering we settle on a GARCH(1,1) model for the data set.




Figure 8: The ACF and PACF of the heavy-tailed simulation.

The histogram of the residuals of the GARCH(1,1) fit reveals a leptokurtic structure (Figure 9(a)). The GARCH(1,1) residuals also result in rejection of the null hypothesis of the Box-Ljung test that they are i.i.d. random variables. And the QQ-plot reinforces the rejection, revealing a substantial departure from the Gaussian, particularly in the tails (Figure 9(b)).

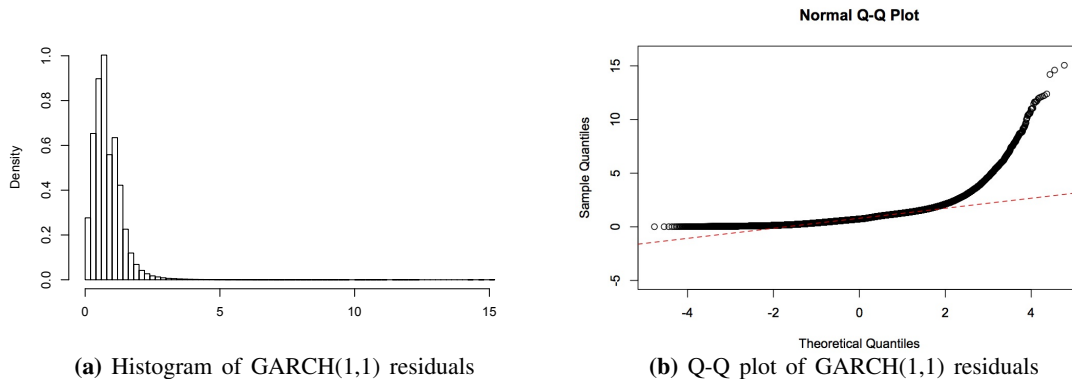


Figure 9: The QQ-Plot of the residuals from the GARCH(1,1) fit to the heavy-tailed simulation.

The above analysis is admittedly rudimentary, but the intention is to illustrate the difficulty in applying some of the standard, readily available techniques to time series generated by agent-based financial models. One could certainly continue to manipulate the parameters of ARMA or GARCH processes to obtain considerably better fits than those presented here. However, this process would not be straightforward, and the researcher would have to rely on various heuristics that ultimately may reduce to subjective judgement about the process. Additionally, these data sets are generated by a known simulation structure where the ARMA and GARCH assumptions were not incorporated directly, and thus their lack of fit is not entirely unexpected.

4.3 Multifractal Results

Visually inspecting the price path from the thin-tailed simulation in Figure 3(b) along with its first difference in Figure 4 reveals a stochastic structure similar to Brownian motion. The erratic nature of the price path indicates a continuous function that is likely too rough to be differentiable, but the first difference indicates the increments of the process may simply be white noise. Given the lack of heavy-tailed behavior and the relative stationarity of the first difference, we formulate the hypothesis that the process is monofractal. We can investigate this hypothesis using MF-DFA, which in the case of a monofractal should reveal a single scaling exponent at all points in time. The value of the scaling exponent corresponds to the Hurst exponent H and gives us insight into the autocorrelation structure of the series. Figure 10(a) shows the results of running MF-DFA on the thin-tailed simulation data set.

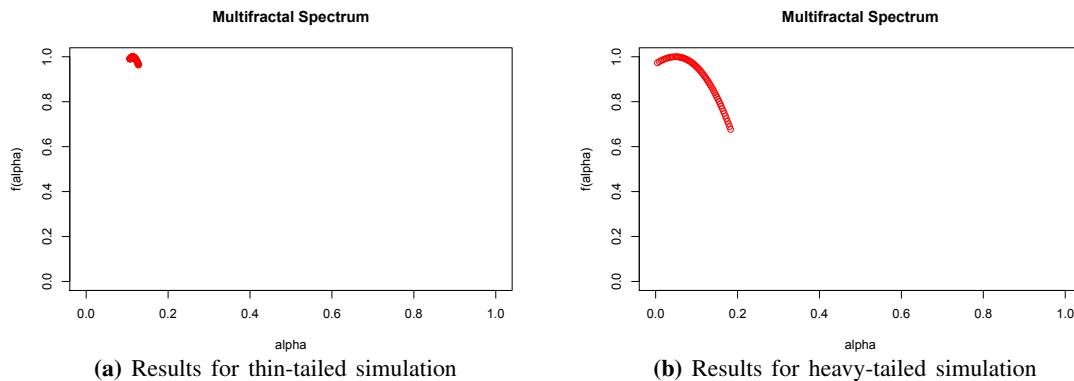


Figure 10: The multifractal spectra for the thin- and heavy-tailed simulations.

The multifractal spectrum confirms the process is monofractal. Essentially the fluctuations within the process all scale in time according to the same generalized Hurst exponent with an approximate value $\alpha = H = 0.11$. In the context of fluctuation analysis, this indicates the process is similar to fractional Brownian motion. The value of $H < 0.5$ indicates the process has short-range dependence and is antipersistent — that is, positive fluctuations are more likely to be followed by negative fluctuations and vice versa. In the context of time series analysis, this procedure represents a more efficient first step in exploratory analysis. By running MF-DFA on the data set first, we are immediately informed that fractional ARMA is likely to be the most appropriate family of models, and short-range dependence along with antipersistence characterize the time series. Multifractal properties are induced by either heavy tails or long-range dependence in a process. Since we deliberately set the distribution for Δ_p to be $N(0,1)$, we conclude that MF-DFA properly detects the thin-tailed nature of the simulation since it reveals a monofractal structure. We also conclude that market discipline, in this case, induces an antipersistence in the price path causing a deviation away from standard Brownian motion, since $H = 0.11 < 0.5$.

Next we apply the same procedure to the heavy-tailed price path. Since the simulation parameters were held exactly the same except for the distribution governing the price offset Δ_p , we might expect to see multifractal properties induced by the heavy-tailed distribution while still maintaining the antipersistent nature we concluded was the result of market structure. The results of running MF-DFA on the heavy-tailed simulation output are shown in Figure 10(b).

The multifractal spectrum of the heavy-tailed simulation output confirms the data set is indeed multifractal, as expected. The range of Hölder exponents is notably broad, but all values of α appear to be less than 0.5. Interestingly, the spectrum peaks at roughly $\alpha = 0.05$ and has a strong concentration of Hölder exponents in the neighborhood of the monofractal produced by the same simulation with a thin-tailed price offset (compare Figures 10(a) and 10(b)). We tested a number of other distributions for the price offset, and confirmed that heavier tails increase the breadth of the multifractal spectrum (and vice versa), but the Hölder exponents remain in the same general neighborhood.

In summary we have shown that zero-intelligence agent-based financial models have the potential to produce monofractal price paths with Hurst exponents similar to fractional Brownian motion, as well as more complex structures with multifractal properties. In our model, the market discipline determined by the order rates and continuous double auction order book structure appeared to induce an antipersistence that was present regardless of which distribution was used for the price offset Δ_p . This is in keeping with conclusions of Farmer, Patelli, and Zovko (2005) and Gode and Sunder (1993) that market structure can produce an aggregate emergent effect that is independent of how irrational the individual participants become.

5 CONCLUSIONS

The zero-intelligence output is an indication that a simple set of rules being followed by 100 participants who are driven by relatively simple random factors can produce a complex time series exhibiting volatility clustering and heavy-tailed marginals. These admittedly broad results are in accordance with previously published accounts of zero-intelligence models, to which we add the conclusion that the time series produced can be multifractal. Gode and Sunder (1993) suggested that Adam Smith's invisible hand may in fact result from the market discipline introduced by procurement processes such as the double auction order book. Farmer, Patelli, and Zovko (2005) also concluded that market efficiency could (at least in theory) emerge on the macrolevel regardless of how irrationally the market participants were behaving. Our experiment with the zero-intelligence model appears to reinforce these conclusions. What is interesting to note is the apparent ability of MF-DFA to differentiate between properties within the time series that were induced by market discipline versus those induced by agent behavior. In our experiment, we altered the agent behavior from placing relatively reasonable bid and ask orders to placing extremely unreasonable bid and ask orders. However, both time series exhibited antipersistence with the dominant (i.e. expected) Hölder exponent in the range of $\alpha \in (0.05, 0.11)$. Since the market structure remained unchanged and only the agent behavior was altered, a reasonable hypothesis for investigation is that market structure is revealed by the dominant Hölder exponent in the output while the degree of multifractality (i.e. the breadth of the multifractal spectrum) is determined by agent behavior. The zero-intelligence model combined with multifractal analysis offers a number of avenues for future research. For example, the quantity of traders and the ratio of patient to impatient traders could be altered to analyze the effects they have on the price paths and the multifractal spectrum.

Missing from this analysis are the calibration of rates to an actual asset and the open question of how to induce persistent (i.e. $\alpha > 0.5$) behavior in the simulated price path. Figure 11 shows the multifractal spectrum from the heavy-tailed simulation compared with the multifractal spectrum from the GE TAQ data set. The GE data set exhibits both antipersistence in the large fluctuations and persistence in the small fluctuations. But its dominant Hölder exponent is roughly $\alpha = 0.48$, while the breadth of exponents exhibited by the series is wider than that of the heavy-tailed simulation. The added complexity of long memory appears to combine with heavy tails in the GE data to produce a wider spectrum than heavy tails alone. Our analysis shows that advanced time series analysis is a critical step in validating agent-based financial markets against empirical data, and yet traditional autoregressive techniques fall short of adequately characterizing real financial data sets. We contend that multifractal analysis reveals the true complexity driving financial markets, and as such has enormous potential for validating agent-based simulations attempting to capture the true nature of financial systems.

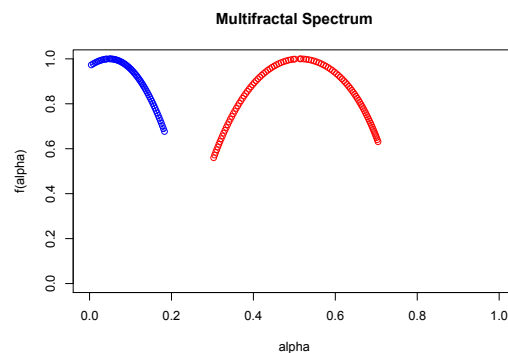


Figure 11: The multifractal spectrum from the GE TAQ data set (blue) compared with that of the heavy-tailed simulation (red).

REFERENCES

- Addison, P. S. 1997. *Fractals and Chaos: An Illustrated Course*. Bristol, UK: Institute of Physics Publishing.
- Calvet, L., and A. Fisher. 2002. "Multifractality in Asset Returns: Theory and Evidence." *The Review of Economics Statistics* 84 (3): 381–406.
- Coffey, W. T., and Y. P. Kalmykov. 2006. *Advances in Chemical Physics: Fractals, Diffusion and Relaxation in Disordered Complex Systems*. Hoboken, NJ: Wiley.
- Engle, R. F. 1982. "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation." *Econometrica* 50 (4): 987–1007.
- Farmer, J. D., P. Patelli, and I. I. Zovko. 2005. "The Predictive Power of Zero Intelligence in Financial Markets." *Proceedings of the National Academy of Sciences of the United States of America* 102 (6): 2254–2259.
- Gode, D. K., and S. Sunder. 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.
- Graneroa, M., J. Segoviab, and J. G. Perez. 2008. "Some Comments on Hurst Exponent and the Long Memory Processes on Capital Markets." *Physica A* 387:5543–5551.
- Kantelhardt, J. W., S. A. Zschiegner, E. Koscielny-Bunde, S. Havlin, A. Bunde, and H. E. Stanley. 2002. "Multifractal Detrended Fluctuation Analysis of Nonstationary Time Series." *Physica A* 316:87–114.
- Mandelbrot, B. B. 1977. *The Fractal Geometry of Nature*. Updated and augmented. New York: W. H. Freeman and Company.
- Mandelbrot, B. B., and J. W. van Ness. 1968. "Fractional Brownian Motions, Fractional Noises and Applications." *SIAM Review* 10 (4): 422–437.
- North, M. J., and C. M. Macal. 2007. *Managing Business Complexity: Discovering Strategic Solutions with Agent-Based Modeling and Simulation*. New York: Oxford University Press.
- Tesfatsion, L. 2006. "Agent-Based Computational Economics: A Constructive Approach to Economic Theory." In *Handbook of Computational Economics: Vol. 2, Agent-Based Computational Economics*, edited by L. Tesfatsion and K. L. Judd, 831–880. San Diego: Elsevier Science & Technology Books.
- Thompson, J. R. 2013. *Analysis of Market Returns Using Multifractal Time Series and Agent-Based Simulation*. Ph.D. thesis, Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh NC. <http://www.lib.ncsu.edu/resolver/1840.16/8781>.

AUTHOR BIOGRAPHIES

JAMES R. THOMPSON recently completed the Ph.D. degree in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. His research interests include stochastic modeling techniques including agent-based simulations, large-scale discrete-event simulations, and systems dynamics as applied to problems of policy-making and decision-analysis in the context of government, economics, and healthcare. His email address is jrthomp5@ncsu.edu and his web page is www4.ncsu.edu/~jrthomp5.

JAMES R. WILSON is a professor in the Edward P. Fitts Department of Industrial and Systems Engineering at North Carolina State University. He is a member of ACM, ASA, ASEE and SCS; and he is a Fellow of IIE and INFORMS. His e-mail address is jwilson@ncsu.edu, and his web page is www.ise.ncsu.edu/jwilson.