Proceedings of the 2014 Winter Simulation Conference
A. Tolk, S. Y. Diallo, I. O. Ryzhov, L. Yilmaz, S. Buckley, and J. A. Miller, eds.

MULTIFRACTAL TIME SERIES ANALYSIS OF POSITIVE-INTELLIGENCE AGENT-BASED SIMULATIONS OF FINANCIAL MARKETS

James R. Thompson

MITRE Corporation
7515 Colshire Dr.
McLean, VA 22102, USA

James R. Wilson

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

ABSTRACT

To analyze the impact of intelligent traders with differing fundamental motivations on agent-based simulations of financial markets, we extend the classical zero-intelligence model of financial markets to a positive-intelligence model using the MASON agent-based modeling framework. We exploit multifractal detrended fluctuation analysis (MF-DFA) to analyze the series of stock prices generated by the positive-intelligence simulation. We study the changes in this output process as analyzed by MF-DFA when altering the mix of agents with competing market philosophies; and we compare and contrast the results of fitting conventional time series models to such output processes with the results of applying MF-DFA to the same processes.

1 INTRODUCTION

The typical economic benchmark for competitive market systems is a Walrasian equilibrium which, by the nature of its assumptions, must contain a market-clearing agent (Tesfatsion 2006). This theoretical agent ensures that demand and prices are optimal under expected value, and that the aggregate supply is greater than or equal to aggregate demand. In the real market, this centralized agent does not exist. Instead the pricing process is governed by a decentralized collection of procurement processes created by market participants who randomly enter and exit the marketplace. Tesfatsion (2006) notes that the challenge to any modeler trying to capture such a tumultuous environment involves, at a minimum, the issues of asymmetric information, strategic interaction, collusion, coordination failure, and the impact of market protocols enforced differently by different governments in a global economy.

Agent-based simulations present us with new possibilities for modeling such environments. Since agent-based models are known to produce emergent effects that were never explicitly programmed by the modeler, some studies test the hypothesis that a Walrasian clearing agent emerges either from market structure, agent behavior, or a combination of the two. Gode and Sunder (1993) constructed a zero-intelligence model to determine if a double auction order book could produce "allocative efficiency" similar to Walrasian equilibrium. The agents were considered to have zero-intelligence because they placed buy and sell orders according to a random process rather than making informed decisions. Gode and Sunder concluded that market discipline (i.e., the set of rules associated with buying and selling) plays a significant role in producing a Walrasian clearing effect.

A distinctive trait of many assets is that their price paths exhibit volatility clustering, meaning periods of relative calm alternate with periods of relative turmoil (LeBaron 2006). This implies that the market is at least partly predictable, and many *technical traders* attempt to exploit this property through econometric analysis of price charts. In contrast to technical traders are the *fundamentalists*, who believe there is a fundamental value associated with the firm, and the asset price will tend toward this fundamental value. These two categories are by no means comprehensive or mutually exclusive. People change their minds and beliefs. They experience feelings of panic and elation. And their economic situations change, not only through changes in their own income and consumption levels, but also because of world events that impact the economy in unpredictable ways. Even within the communities of technical traders and fundamentalists, there are disagreements over which method of economic analysis to use or what constitutes the fundamental value of the firm. The possibilities are innumerable, and again we are faced with the question of which combination of these microlevel behaviors and beliefs actually impact the market and which are obscured or canceled out on the macrolevel.

North and Macal (2007) advocate an approach to agent-based modeling that starts with protoagents which can be easily embellished in an incremental fashion. In keeping with this approach, in an earlier paper we presented an analysis of a zero-intelligence model (Thompson and Wilson 2013). In this paper we extend the zero-intelligence model to incorporate *positive-intelligence* traders that fall into two groups: fundamental traders we call *Grahamists*, and technical traders we call *Chartists*. The objective of the positive-intelligence model is to progress toward reality in an incremental way while maintaining a simplicity that lends itself to systematic experimentation. The economic debate on market behavior is ongoing and unlikely to be resolved by the use of simple models. However, insight can be gleaned from systematic exploration of tractable models. Our intention is to analyze the impact of added intelligence on the multifractal properties of the simulated price paths, and to compare and contrast traditional methods of time series analysis with multifractal analysis on the outputs of agent-based financial markets.

In this article, we exploit MF-DFA to analyze the output of agent-based simulations of financial markets—specifically, the simulation-generated price paths of given assets. In Section 2 we introduce the concept of fractals and multifractals in the context of time series analysis. We present a high-level overview of the multifractal spectrum, and summarize MF-DFA as a method for estimating the multifractal spectrum from empirical data. In Section 3 we extend the zero-intelligence model to accommodate positive-intelligence traders. We describe the enhancements we made to the double auction order book and give an in-depth description of how we incorporated market philosophies into our trading agents. In Section 4 we describe our different experiments with the positive-intelligence model; and using MF-DFA, we compare the different methods of analyzing financial time series. We close in Section 5 with a brief discussion of our main conclusions. The slides for the oral presentation of this article are available online via www.ise.ncsu.edu/jwilson/wsc14mfdfa.pdf.

2 MULTIFRACTAL DETRENDED FLUCTUATION ANALYSIS

As outlined in Thompson and Wilson (2013) and elaborated in Thompson and Wilson (2014), for a self-similar fractal object, each small part of the object resembles the whole object in some essential respect. For example, a Brownian motion process $\{X(t) : t \geq 0\}$ is self-similar because the increment $X(t+s)-X(t)$ of the process over the fixed time interval $[t, t+s]$ is a suitably rescaled probabilistic replica of the increment $X(t+\eta s)-X(t)$ over the much smaller time interval $[t, t+\eta s]$ when $0 < \eta \ll 1$. In particular, $X(t+s)-X(t)$ is normal with mean 0 and variance s while $X(t+\eta s)-X(t)$ is normal with mean zero and variance ηs; and when $\eta \gg 1$, a similar relationship holds. The Brownian motion process $\{X(t) : t \geq 0\}$ has sample paths that are continuous, nondifferentiable, and have Hölder exponent $\alpha(t) = 0.5$ at every time $t \geq 0$. Roughly speaking, the latter property means that for each $t \geq 0$, the increment $X(t+s)-X(t)$ of the process is with high probability of the order of $\alpha(t) = |s|^{0.5}$ as the time-interval length $|s| \to 0$; and this result is easily verified because $X(t+s)-X(t)$ is normally distributed with mean zero and standard deviation $|s|^{0.5}$ for all $t \geq 0$ and $|s| \leq t$.

Examination of self-similar data sets in the context of time series analysis shows that the Hurst exponent H characterizes the asymptotic behavior of the autocorrelation function (ACF) of the time series (Mandelbrot and van Ness 1968). For example in the case of Brownian motion, we have $H = \alpha = 0.5$. Values of H in the interval $(0.5, 1.0)$ lead to positive autocorrelations that decay too slowly for the sum of autocorrelations over all lags to be finite. Processes with $0.5 < H < 1$ are said to exhibit long-range dependence (long memory); and processes with $0 \le H \le 0.5$ are said to exhibit short-range dependence (short memory).

A process with $0 < H < 0.5$ exhibits *antipersistence* in its sample paths, which means that a positive increment (increase) in the process is more likely to be followed by a negative increment (decrease) in the next nonoverlapping time interval and vice versa; and this tendency of the process to turn back on itself results in sample paths with a very rough structure. When $0.5 < H < 1$, the process exhibits *persistence*, which means that successive nonoverlapping increments in the process are more likely to have the same sign; and smoother sample paths result from this tendency of the process to persist in its current direction of movement. Therefore in a self-similar process with Hurst exponent H (that is, a "monofractal" process), H quantifies not only the asymptotic behavior of the ACF but also the inherent roughness of the sample paths of the process. Brownian motion exhibits neither antipersistence nor persistence because its nonoverlapping increments are independent; moreover, its sample paths are nonsmooth in the sense of being continuous but not differentiable. It can be shown that the Hausdorff dimension D of the sample path of a monofractal Gaussian process (i.e., Brownian motion or fractional Brownian motion) is related to the Hurst exponent of the underlying process by the relation $H = 2 - D$. For example, Brownian motion has $H = \alpha = 0.5$ so that every sample path of Brownian motion has Hausdorff dimension $D = 1.5$.

A multifractal process $\{X(t) : t \in [0,T]\}$ is defined in a similar way to monofractals, except that the Hurst exponent H is no longer a single scalar value, but rather a function that ultimately depends on $\{\alpha(t) : t \in [0,T]\}$, the set of Hölder exponents characterizing the local behavior of the process over its finite time horizon. If $\mathscr{Q}$ is a certain neighborhood of zero that contains the unit interval $[0,1]$ and $q \in \mathscr{Q}$, then the generalized Hurst exponent $h(q)$ is defined in terms of the qth absolute moment of $X(t)$ for each $t \in [0,T]$. In contrast to the local Hölder exponents, $h(q)$ is a global characteristic of the process over the entire time horizon $[0,T]$. Derived from the generalized Hurst exponents $\{h(q) : q \in \mathscr{Q}\}$, the multifractal spectrum $f(\alpha)$ for $\alpha \ge 0$ provides a concise description not only of the general arrangement of the local Hölder exponents over the time horizon of the process, but also of the way in which the sample paths of the process exhibit short- or long-range dependence. For a more complete discussion of multifractals, see §§1–2 of Thompson and Wilson (2014).

One way to think of a multifractal process $\{X(t) : t \in [0,T]\}$ is as the amalgamation of an infinite number of monofractal subprocesses, each characterized by a single Hölder exponent α. However, these monofractal processes are interwoven throughout the time horizon $[0,T]$ such that the set of time points associated with any one monofractal process constitutes a fractal set. The function $f(\alpha)$ describes the following key properties of a multifractal time series:

- The set of Hölder exponents $\{\alpha(t) : t \in [0,T]\}$ specify how the underlying stochastic process $\{X(t) : t \in [0,T]\}$ fluctuates as we examine its increments computed from nonoverlapping time intervals whose common length is systematically varied over a broad range of values.
- For a particular nonnegative value α_0 of the Hölder exponent, the corresponding value $f(\alpha_0)$ of the multifractal spectrum is the Hausdorff dimension of the subset of time points $t \in [0,T]$ at which the stochastic process $\{X(t) : t \in [0,T]\}$ has its Hölder exponent $\alpha(t) = \alpha_0$.

Although $f(\alpha)$ is not a proper probability density function for $\alpha \ge 0$, in its "renormalized" form for each fixed α_0 the associated function value $f(\alpha_0)$ represents the general arrangement of time points at which the multifractal process $\{X(t) : t \in [0,T]\}$ has the specific value α_0 for its Hölder exponent; see Section IV.A of Calvet and Fisher (2002).

The multifractal spectrum $f(\alpha)$ is concave with negative second derivative for $\alpha \ge 0$, achieving the value of 1 at its unique global maximum. If a stochastic process is monofractal, then it has a single Hölder exponent that coincides with its Hurst exponent H and describes how its increments behave locally at all time points; and therefore the multifractal spectrum of a monofractal process with Hurst exponent H is given by

$$f(\alpha)=\begin{cases}1, & \text{if } \alpha=H,\\ 0, & \text{if } \alpha\neq H \text{ and } \alpha\geq 0.\end{cases} \tag{1}$$

These various interpretations are collectively referred to as the multifractal formalism, and they provide the basis for our intuition about the multifractal spectrum. They also lead to methods for estimating the multifractal spectrum from a given time series.

2.1 The MF-DFA Algorithm

Among the most effective methods for estimating $f(\alpha)$ 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-magnitude fluctuations in the data will be accentuated and the small-magnitude fluctuations will be diminished. Conversely, when q is negative with a large magnitude, the small-magnitude fluctuations will be accentuated and the large-magnitude fluctuations will be diminished. Therefore long memory in fluctuations of different magnitudes (e.g., volatility clustering) can be detected.

The basic MF-DFA algorithm involves the following steps: (i) partitioning a time series of length N into nonoverlapping subseries (segments) of length s (where $20\leq s\leq N/10$); (ii) fitting a polynomial to the observations within each segment; and (iii) computing the qth root of the average taken over all segments of the qth power of each residual standard deviation estimated within a segment in step (ii). This process is repeated as s is systematically varied over a broad range of values satisfying $20\leq s\leq N/10$; and for a fixed value of q, the results of step (iii) for each segment length s are plotted against s on a doubly logarithmic scale. This is analogous to analyzing how the increments of a monofractal process depend on the time-interval length to derive the conventional Hölder (and Hurst) exponent. For each selected value of q, a linear function of s is fitted to the resulting doubly logarithmic plot; and the slope of the fitted linear function is our estimate of the generalized Hurst exponent $h(q)$. From our estimates of $h(q)$ for $q\in\mathscr{Q}$, we can estimate the associated scaling function $\tau(q)$ for $q\in\mathscr{Q}$; and then by taking the Legendre transform of the scaling function $\tau(q)$, we obtain estimates of α and $f(\alpha)$. For a detailed algorithmic statement of the MF-DFA procedure and our robust, computationally efficient method for implementing that procedure, see §§2.4–2.5 of Thompson (2013) or §3 of Thompson and Wilson (2014).

3 THE POSITIVE-INTELLIGENCE MODEL

In this section we extend the zero-intelligence double auction order book model in an effort to progress toward a more realistic model of how markets operate. A significant obstacle in this endeavor is establishing a benchmark for that reality. The mechanics of the order book can certainly be enhanced to allow orders of more than one share and to enforce market capitalization by conserving shares. Most exchanges operate in a double auction order book format, with well-defined rules for trading and market clearing. But accurately defining the realistic behavior of market participants is the subject of an ongoing debate. Our zero-intelligence model exhibited both monofractal antipersistent price paths and multifractal antipersistent price paths (Thompson and Wilson 2013). To induce multifractal properties, the zero-intelligence agents had to utilize an unrealistic distribution for the price offsets. Our objective is to induce multifractal properties through more reasonable agent behavior. Additionally the analysis by Thompson and Wilson (2013) of high-frequency trades of GE stock yielded multifractal spectra that exhibited both antipersistence and persistence in price fluctuations on different time scales. This observation motivated our investigating the question of whether positive-intelligence traders can induce persistence in the market while still interacting with a double auction order book in a plausible way. Our positive-intelligence market simulation was constructed using the MASON agent-based modeling framework (Luke 2005).

3.1 The Order Book

The positive-intelligence financial market model is an extension of the zero-intelligence model of Thompson and Wilson (2013). The market structure is again represented by a double auction order book, meaning agents can place orders to either buy or sell shares of a single asset. The order book accepts limit orders that do not immediately cross the spread and are therefore placed in a best-price-first queue. In the case of bid limit orders, the best price is the highest bid; and conversely in the case of ask limit orders, the best price is the lowest ask. Ties are broken by a first come, first served queuing discipline within the bid- and ask-order queues. The order book also accepts market orders that do not specify a price and are intended to execute at the current best price.

We enhanced the logic of the double auction order book in two ways. First, the order book manages the conservation of shares within the market model. In the zero-intelligence model we were solely focused on how the price path of the stock was impacted by the distribution of price offsets and the processes governing placement and expiration of orders. As such, the agents did not keep track of their wealth; and we assumed an infinite supply of shares. In the positive-intelligence model, we require the agents to react to market conditions, which includes the possibilities of excess demand or excess supply. Conserving shares provides a mechanism for causing such conditions. This means bid orders can only execute if there are orders in the ask-order queue, and ask orders can only execute if there are orders in the bid-order queue.

Second, we adjusted the order-generation process to allow agents to place individual orders for more than one share of the asset. As we discuss in greater detail below, the agents now have limited wealth with which to operate, so their entry and exit from the market requires an additional degree of freedom that is facilitated by being able to buy or sell multiple shares at once. The order book logic was enhanced to search the appropriate queue for crossed orders and satisfy incoming orders up to the available quantity of shares. In the case of market orders, there is no specified price. This means one bid order may be satisfied by multiple ask orders to provide the quantity of shares being ordered; similarly, one ask order may be satisfied by multiple bid orders to provide the required order quantity.

Each time a transaction occurs, the double auction order book updates the market price that is visible to all agents. If the transaction is an incoming market order, then the order book calculates a share-weighted average of the transaction to determine the market price. For example, if a market bid order requests three shares and there is one share being offered at \$100.00 and two shares being offered at \$102.00, then the market bid order is fully satisfied and the new market price is $[1(\$100)+2(\$102)]/3 = \$101.33$. Immediately following a market order, the new bid-ask spread may in fact be negative, meaning existing limit orders are crossed and can be satisfied. The order book handles this situation by taking the midpoint between the two unit share prices. Essentially, the bidding agent is getting her shares at a slightly cheaper price than she was willing to pay, and the asking agent is selling his shares at a slightly higher price than he was willing to offer. The midpoint means their respective gains are as fair as possible.

3.2 The Agents

The most significant difference in the positive-intelligence model is the behavior of the agents. In the zero-intelligence model there were patient traders placing only limit orders and impatient traders placing only market orders. Beyond being able to query the order queues for the best price and being able to generate random variables for setting prices and order interarrival times, the agents had no regard for market conditions. In the positive-intelligence model, the agents must make decisions in the face of scarce resources. Each agent is instantiated with a certain amount of wealth, split between a risk-free asset, called the bond, and an initial stipend of shares. These values are determined randomly at simulation clock time $t = 0$, but the maximum values can be altered by the user. As previously mentioned, the order book conserves shares, so the initial instantiation process determines the market capitalization. Agents can immediately sell the shares they possess or attempt to buy more shares, depending on their internal logic and decision thresholds. However, they will only be successful if other agents are simultaneously attempting the opposite operation. We chose to have 500 agents classified into two groups: Grahamists and Chartists. The user can specify the percentage of the 500 agents that fall into either category. Agents of either classification have the ability to place either market or limit orders. They can query the order book to determine the current price of the asset and then use that price to determine their current wealth. Agent i earns interest on that agent's bond and possesses a parameter β_i specifying the fraction of the individual's total wealth that agent is not willing to invest in the risky asset. Each agent also has a minimum wealth threshold. If an agent's total wealth drops below that threshold, then the agent makes no attempt to buy or sell until the individual's wealth increases above the threshold.

3.2.1 Grahamists

Grahamists, as their name implies, follow the practical rule set forth by Graham ([1949] 2005). Specifically, each investor should choose a fraction β of the individual's wealth that the investor is not willing to risk. This fraction should be invested in low-risk (and therefore lower-yielding) bonds. The remaining fraction $1-\beta$ should be invested in firms that the investor believes possess value. The task of the investor is then to maintain this division of wealth. If the share price increases, then the investor sells enough shares to rebalance the chosen bond-to-asset split. If the share price decreases, then the investor uses some of the bond money to purchase more shares. The net result is buying low and selling high, while maintaining a steady long-term growth in wealth. This strategy is based on the belief that markets are efficient, but in the short run irrational investors can cause a mispricing of a given asset that is soon corrected by fundamentals. There are a number of criticisms of this belief, but they are well summarized by a quote of John Maynard Keynes: "Markets can remain irrational a lot longer than you and I can remain solvent" (Schilling 1993). The obvious implication is that if investors continually shift their bond holdings into a decreasing market, they may ultimately exhaust all their wealth before the market fundamentals intervene to correct the mispricing.

Upon instantiation at simulation clock time $t=0$, the Grahamist with agent index i receives a random amount of wealth up to a threshold set by the user, along with the main decision parameter β_i. For Grahamist agent i, the parameter β_i is randomly sampled in the interval $(0,1)$; and that individual is immediately allocated enough shares to account for the fraction $1-\beta_i$ of the agent's current wealth at the initial share price. Thus, at simulation clock time zero, Grahamist agent i holds the fraction β_i of that agent's total wealth in the risk-free bond and the fraction $1-\beta_i$ of the agent's total wealth in shares of the risky asset. Grahamist agent i then schedules a stock-price query at a simulation clock time in the near future with rate λ_i so that the time between successive queries of the asset price is randomly sampled from the exponential distribution with mean $1/\lambda_i$. The query rate λ_i of Grahamist agent i is randomly sampled from the uniform distribution on the half-open interval $(0,2]$ to ensure heterogeneity in the behavior patterns of Grahamist agents. However, the user can adjust the maximum value of the half-open interval for different runs of the simulation.

Each time a Grahamist agent queries the market to determine the share price, the agent uses that information along with the risk-free interest rate to update the agent's wealth. The wealth $W_{t,i}$ of Grahamist agent i at simulation clock time t is determined by the relation

$$W_{t,i} = P_t S_{t,i} + B_{t_i^*,i}\left(1+r_f\right)^{t-t_i^*}, \tag{2}$$

where: (a) P_t is the current market price; (b) $S_{t,i}$ is the current quantity of shares held by agent i at simulation clock time t; and (c) $B_{t_i^*,i}$ is the amount held by agent i in risk-free bonds with the fixed daily interest rate r_f since the last simulation clock time t_i^* (expressed in days) at which agent i queried the market for the current stock price. The daily interest rate can be adjusted by the user before starting the simulation, and by default it is set to $r_f = 0.03/365$. Note the compounding factor $\left(1+r_f\right)^{t-t_i^*}$ accounts for the elapsed simulation clock time (in days) that the amount $B_{t_i^*,i}$ was held in bonds by Grahamist agent i.

Once Grahamist agent i has queried the double auction order book for the current market price, that agent checks for deviations from the desired bond-to-asset split as specified by the agent's assigned value of

β_i. If the market price has increased, then Grahamist agent i attempts to sell enough shares to reestablish the desired bond-to-asset split. Conversely if the market price has decreased, then Grahamist agent i attempts to buy enough shares to reestablish the β_i split. The first attempt to reestablish this split is done with a market order. Because of the conservation of shares enforced by the double auction order book, this attempt by Grahamist agent i will be unsuccessful if no other agent has placed a limit order of the opposite operation. If such an attempt fails, then Grahamist agent i offsets the current market price P_t by an amount U randomly sampled from a folded normal distribution so that $U = |Z|$, where $Z \sim \text{Normal}(0, \sigma^2)$, and σ^2 can be adjusted by the simulation user and is set to $\sigma^2 = 1$ by default. If agent i is placing a bid, then the agent's new bid price is $P_t + U$; and if agent i is placing an ask, then the agent's new ask price is $\max\{P_t - U, 0\}$. Agent i then recalculates the split based on this new price and places a limit order with this new quantity and price.

It is important to note that there are 500 total agents, all attempting to buy and sell through a single order book. To handle this situation in Java, a certain degree of synchronization is required. Synchronization essentially locks a resource (or operation) when it is being used by an agent. As soon as the current agent relinquishes the resource, that resource is available to be employed by the next agent. A Grahamist agent queries the double auction order book for the current price, makes a decision, and finally places a market order. The market order does not specify a price, but it does specify a quantity, which is a function of the price that the Grahamist agent observed when querying the market. Because of synchronization, by the time the agent's market order is able to access the double auction order book to place the order, the price may have changed substantially from the original price used to calculate the Grahamist agent's split. Given enough agents in the market, this process could place the Grahamists in a constant state of flux trying to reestablish their initial split. We dampened this effect by creating a tolerance level for the β_i-split that can be adjusted by the user and is set to 0.005 by default. That is, the Grahamists only attempt to reestablish their split when their bond fraction of wealth differs from β_i by more than the tolerance level.

3.2.2 Chartists

While Grahamists represent our fundamental traders, Chartists represent the technical analysts who depend on econometric forecasts of market movements to make their decisions. Like a Grahamist agent, Chartist agent i receives an allocation of wealth at instantiation that is immediately split between the risk-free asset and the risky asset according to the agent's assigned parameter β_i. However, Chartists do not seek to maintain a fixed bond-to-asset ratio. Instead as elaborated below, the parameter β_i for Chartist agent i serves as a measure of risk aversion in that $1 - \beta_i$ helps determine the quantity of shares that Chartist agent i is willing to risk on the next stock-price forecast, which always coincides with the agent's next stock-price query.

Chartist agent i also possesses a parameter ψ_i that determines the agent's forecast of the next stock price. For simplicity, all Chartists utilize exponential smoothing as their forecast methodology. For Chartist agent i, let ω_i denote the index of the agent's next stock-price query (where $\omega_i = 1, 2, \ldots$), which will be made at simulation clock time $t(\omega_i, i)$ (where $t(0, i) \equiv 0$ for every Chartist agent i) and will yield the actual share price $P_{t(\omega_i, i)}$ at that point in simulation clock time. Immediately after stock-price query $\omega_i - 1$ occurring at simulation clock time $t(\omega_i - 1, i)$, Chartist agent i computes the forecast $\widehat{P}_{\omega_i, i}$ of the stock price that will be observed at the time of the agent's next stock-price query according to the exponential smoothing formula

$$\widehat{P}_{\omega_i, i} = \psi_i P_{t(\omega_i - 1, i)} + (1 - \psi_i)\widehat{P}_{\omega_i - 1, i}\,, \tag{3}$$

where: (a) ψ_i is the smoothing constant for Chartist agent i; (b) $P_{t(\omega_i - 1, i)}$ is the current stock price at simulation clock time $t(\omega_i - 1, i)$; and (c) $\widehat{P}_{\omega_i - 1, i}$ represents the last forecast made at simulation clock time $t(\omega_i - 2, i)$ of the actual stock price observed at time $t(\omega_i - 1, i)$. Note that for each Chartist agent i, the initial condition for the recursion (3) is established by offsetting the initial stock price P_0 at simulation clock time $t = 0$ with a random sample $Z_i \sim \text{Normal}(0, 1)$ so that we have $\widehat{P}_{0,i} = \max\{P_0 + Z_i, 0\}$. The

Chartists schedule themselves to update their forecasts and make decisions about market actions according to a homogeneous Poisson process with market-query rate λ_i. Therefore the amount of simulation clock time that elapses between stock-price queries $\omega_i - 1$ and ω_i for Chartist agent i is exponentially distributed with mean $1/\lambda_i$.

To maintain a degree of heterogeneity within the Chartist community, the smoothing constant ψ_i is set randomly in the interval $(0, 0.9]$ for Chartist agent i. When Chartist agent i advances to stock-price query ω_i at simulation clock time $t(\omega_i, i)$, that agent first observes the current market price $P_{t(\omega_i,i)}$ and then updates the agent's wealth according to Equation (2) in which $t = t(\omega_i, i)$ and $t_i^* = t(\omega_i - 1, i)$. Chartist agent i then updates the relevant query index by the assignment $\omega_i \leftarrow \omega_i + 1$ and computes the agent's forecast of the stock price at the time of the next query according to Equation (3).

The goal of the Chartist is to make money on market movements rather than long-term growth. So if Chartist agent i forecasts an increase in price, then the agent places a market order to buy shares up to the fraction $1 - \beta_i$ of the agent's current wealth; and simultaneously Chartist agent i places a limit order to sell that same quantity of shares at the agent's current forecasted price. Assuming both orders are successfully executed over some interval of simulation clock time required for the limit order to be crossed, Chartist agent i will realize a profit of $\big[\widehat{P}_{\omega_i,i} - P_{t(\omega_i-1,i)}\big] S_{t(\omega_i-1,i),i}$, where $S_{t(\omega_i-1,i),i}$ is the quantity of shares held by that agent immediately after the attempted execution of the agent's market order placed at the current time $t(\omega_i - 1, i)$. If Chartist agent i is unable to execute the market order to purchase shares, then the agent still attempts to place a limit order to sell either the original quantity of shares or the agent's current holding of shares, whichever is less. The structure of our market does not allow short selling, so the Chartists are limited by their buy market order and their current holdings.

If Chartist agent i forecasts a decrease in price, then the agent places a market order to sell shares up to the fraction $1 - \beta_i$ of the agent's current wealth or the number of shares currently held, whichever is less. Chartist agent i simultaneously places a limit order to buy back those shares at the forecasted lower price $\widehat{P}_{\omega_i,i}$. Regardless of whether Chartist agent i is able to execute the sell market order, the agent proceeds with the buy limit order. The idea is that purchasing at the lower price will get Chartist agent i back into the market and in a position to profit on future changes in the stock price.

All limit orders are limited by the agent's wealth in that a Chartist agent cannot place an order to buy at a value greater than the agent's current wealth. This is not to say that agents cannot go bankrupt. Because of synchronization, they may believe they can afford a market order at the current price, but market orders ahead of theirs may shift that price before they are able to execute. If the shift is large enough, then the agent will buy an order that exceeds their current wealth, causing that agent to exit the market altogether.

3.3 Experiments

Our objective was to analyze the changes in the multifractal properties caused by adding positive intelligence to the agents, while simultaneously incorporating only plausible behavioral traits. In the zero-intelligence model, the nature of the double auction order book appeared to create only antipersistent behavior, as large market orders created a deeply negative spread that gave rise to very low bids and very high asks being satisfied in rapid succession. Since real markets also exhibit persistent behavior, we tested whether a double auction order book combined with intelligent traders could produce this effect. We also analyzed the impact of different market compositions on the multifractal spectrum. There is still the open question of which microlevel behaviors interact to produce notable macrolevel effects versus those that are obscured on the aggregate level. By systematically increasing the proportion of Grahamists in the model, we analyzed the changes to the price of the stock and the multifractal properties induced by shifting the predominant thinking in the market from Chartist to Grahamist. For all the experiments, we conducted standard exploratory time series analysis to confirm the claims by other researchers that agent-based models can replicate the stylized facts of real financial markets (LeBaron 2001, LeBaron 2012, Farmer, Patelli, and Zovko 2005). To these analyses we added multifractal analysis to confirm that generalized Hurst exponents and self-similarity are also properties that can emerge from agent-based models of financial markets.

4 RESULTS

4.1 Traditional Time Series Analysis

To illustrate the difficultly in analyzing financial time series, we examined the simulation output for the case of 80% Grahamists in greater detail. We started with the first difference of the price and the first difference of the log of the price as shown in Figure 1. We concluded that both time series exhibited volatility clustering and that taking the log returns failed to dampen this effect. We also analyzed the autocorrelation and partial autocorrelation functions of the original series and the log-transformed series in an attempt to detect autoregressive moving average (ARMA) behavior in the mean (see Figure 2).

The exploratory analysis indicated the volatility in the time series was nonconstant and might be well characterized by a GARCH model. The autocorrelation structure also indicated that the mean was not stationary and might follow an autoregressive process. Therefore we attempted to fit combined ARMA(p,q)+GARCH(1,1) models to both the first difference of the price and the first difference of the log price. However, the fGarch package in R encountered a singularity error when trying to obtain fits to the differenced time series (Wuertz and Chalabi 2013). Ultimately we were able to obtain a fit to the log price. Figure 3 shows a QQ-plot and a histogram of the residuals from the ARMA(2,0)+GARCH(1,1) model fitted to the log price generated by the first replication of the market simulation with 80% Grahamists. Note the leptokurtic nature of the histogram and the drastic departure of the tails from the normal. The visual inspection of these plots indicated that the model was a poor fit, and that conclusion was supported by Box-Ljung-Pierce portmanteau lack-of-fit test. The p-value from the latter test was less than 2.2×10^{-16}, indicating the ARMA(2,0)+GARCH(1,1) model was an inadequate description of the time series.

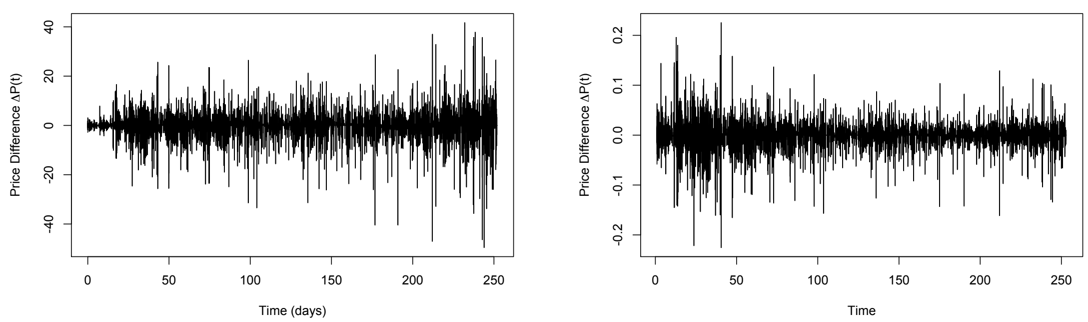

(a) Price difference for 80% Grahamists **(b)** Difference of the log price for 80% Grahamists

Figure 1: Results generated by the first run of the market simulation with 80% Grahamists.

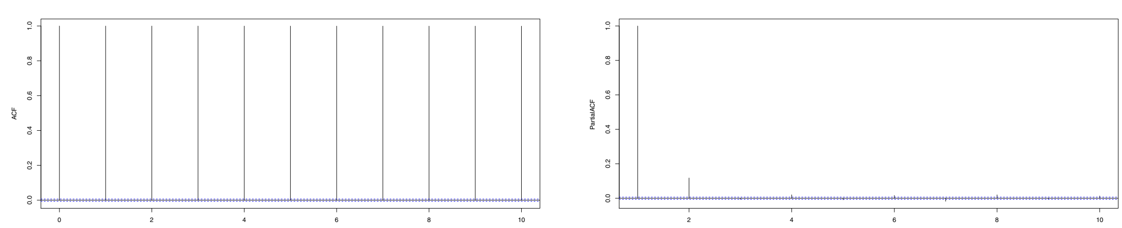

(a) Autocorrelation function of log price **(b)** Partial autocorrelation function of log price

Figure 2: The autocorrelation structure of log price time series generated by the first run of the market simulation with 80% Grahamists.

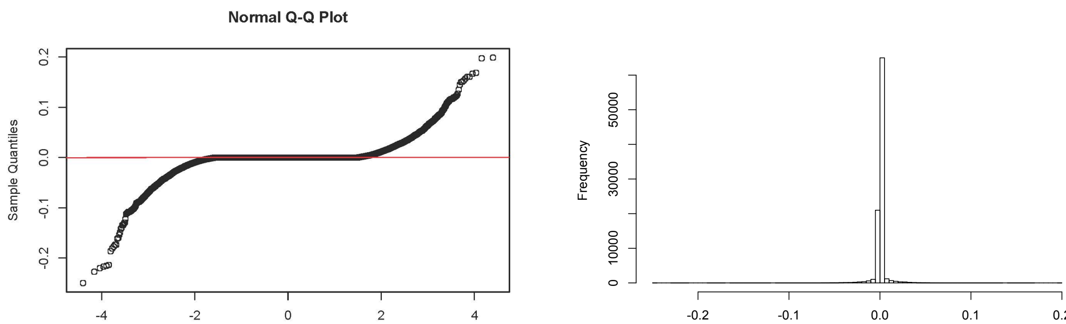

(a) QQ-plot of residuals **(b)** Histogram of residuals

Figure 3: Analysis of residuals from the ARMA(2,0)+GARCH(1,1) model fitted to the log price time series generated by the first run of the market simulation with 80% Grahamists.

4.2 Multifractal Time Series Analysis

To investigate the multifractal properties produced by the positive-intelligence financial market, we performed ten runs of the simulation with 80% Grahamists, starting with different random number seeds on each run. Figure 4 shows the multifractal spectra obtained from the replications. The most central spectrum is shown in red, while the others are in gray for illustration purposes.

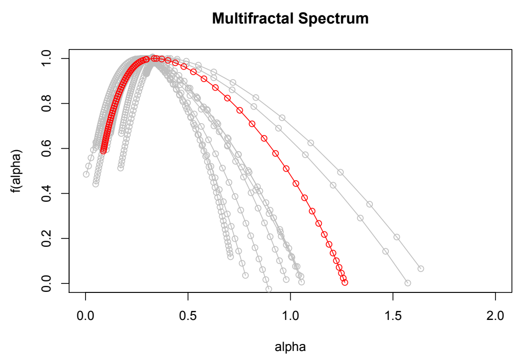

Figure 4: Multifractal spectra for ten runs of 80% Grahamists simulation

Although there was significant variation in the range of Hölder exponents α, there was relatively little variation in the estimate of $\mathrm{E}[\alpha]$ on each replication. For ten replications the average α-value corresponding to $f(\alpha) = 1$ was 0.319 with a 95% confidence interval of $[0.287, 0.350]$. This indicated that there was some aspect of the system's behavior that remained relatively stable and determined the dominant Hölder exponent in the output of each replication of the simulation. As we noted earlier, the multifractal spectrum identifies Hölder exponents that contribute to the overall price path at a vanishing frequency. In the positive-intelligence model, these contributions appeared to be sufficiently large to change the multifractal spectrum from one run to the next, but not large enough to alter the expected value of the Hölder exponent. As detailed in Section 5.6 of Thompson (2013), we observed a similar effect in the GE stock price data by computing the multifractal spectra for the four subseries corresponding to the calendar years 2000, 2001, 2002, and 2003, respectively.

5 CONCLUSIONS

From the results of the positive-intelligence financial market, we concluded that a double auction order book structure coupled with informed agents, making decisions with publicly available information, and using basic philosophies of market dynamics, could produce multifractal time series exhibiting persistence in the price increments. Our analysis also indicated that positive agent intelligence would be required if agent-based models hope to imitate the properties we have observed in empirical financial data. The positive-intelligence model only incorporated thin-tailed distributions for price offsets. Yet, the model produced plausible price paths, some persistence in price increments, and multifractal properties.

We also concluded that without multifractal analysis, we were limited to parametric time series models that, in this case, failed to produce an adequate description of the complex time series. Running MF-DFA on the output of the simulation produced a spectrum that informed us about the autocorrelation structure of the times series, the relative persistence of the increments in the series, and the level of complexity (i.e., monofractal or multifractal) inherent in the time series.

The positive-intelligence model in its current form offers several opportunities for future research. The impact of changing the quantity of agents, step rates of agents, parameter values for each group of agents, and the initial price of the asset, could all be analyzed with little or no modification to the current model. One especially interesting extension of the positive-intelligence model would be to implement an adaptive exponential smoothing technique that effectively improves the forecast accuracy for each Chartist over time; see, for example, Trigg and Leach (1967) and Ekern (1981). However, these essential next steps still fall into the category of computational thought experiments. Another addition to such models would be the shock of external events. Currently the model is a closed system with no outside influences. However, pundits and economists alike claim that certain world events bolster or rattle the confidence of investors. One relatively simple inclusion could be a representation of a Federal Reserve agent that adjusts the interest rates periodically. By allowing the investing agents to borrow capital for investment in the risky market, we could analyze the effects of leverage and the impact of interest rate control by a central bank or other governmental entity.

REFERENCES

Calvet, L., and A. Fisher. 2002. "Multifractality in Asset Returns: Theory and Evidence." *The Review of Economics Statistics* 84 (3): 381–406.

Ekern, S. 1981. "Adaptive Exponential Smoothing Revisited." *Journal of the Operational Research Society* 32 (9): 775–782.

Farmer, J. D., P. Patelli, I. I. Zovko, and K. J. Arrow. 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.

Graham B. (1949) 2003. *The Intelligent Investor*. New York: HarperCollins World.

Kantelhardt, J. W., S. A. Zschiegner, E. Koschielny-Bunde, S. Havlin, A. Bunde, and H. E. Stanley. 2002 "Multifractal Detrended Fluctuation Analysis of Nonstationary Time Series." *Physica A* 316:87–114.

LeBaron, B. 2001. "Empirical Regularities from Interacting Long- and Short-Memory Investors in an Agent-Based Stock Market." *IEEE Transactions on Evolutionary Computation* 5 (5): 442–455.

LeBaron, B. 2006. "Agent-Based Computational Finance." In *Handbook of Computational Economics: Agent Based Computational Economics, Volume II*, edited by L. Tesfatsion and K. L. Judd, 1187–1233. Amsterdam: Elsevier.

LeBaron, B. 2012. "Heterogeneous Gain Learning and the Dynamics of Asset Prices." *Journal of Economic Behavior and Organization* 83 (3): 424–445.

Luke, S., C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan. 2005. "MASON: A Multiagent Simulation Environment." *Simulation* 81 (7): 517–527.

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.

Schilling, A. G. 1993. "Scoreboard." *Forbes* 151 (4): 236.

Tesfatsion, L. 2006. "Agent-Based Computational Economics: A Constructive Approach to Economic Theory." In *Handbook of Computational Economics: Agent Based Computational Economics, Volume II*, edited by L. Tesfatsion and K. L. Judd, 831–880. Amsterdam: Elsevier.

Thompson, J. R. 2013. *Analysis of Market Returns Using Multifractal Time Series and Agent-Based Simulation.* Ph.D. dissertation, Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh North Carolina. Available online via http://www.lib.ncsu.edu/resolver/1840.16/8781 [accessed July 14, 2014].

Thompson, J. R., and J. R. Wilson. 2013. "Multifractal Analysis of Agent-Based Financial Markets." In *Proceedings of the 2013 Winter Simulation Conference*, edited by R. Pasupathy, S.-H. Kim, A. Tolk, R. Hill, and M. E. Kuhl, 1383–1394. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers. Available via http://informs-sim.org/wsc13papers/includes/files/121.pdf [accessed July 14, 2014].

Thompson, J. R., and J. R. Wilson. 2014. "Multifractal Detrended Fluctuation Analysis: Practical Applications to Financial Time Series." Technical Report, Edward P. Fitts Department of Industrial and Systems Engineering, North Carolina State University, Raleigh, North Carolina. Available via http://www.ise.ncsu.edu/jwilson/files/thompson14mfdfa.pdf [accessed July 14, 2014].

Trigg, D. W., and A. G. Leach. 1967. "Exponential Smoothing with an Adaptive Response Rate." *Operational Research* 318 (1): 53–59.

Wuertz, D., and Y. Chalabi. 2013. "Rmetrics—Autoregressive Conditional Heteroskedastic Modelling." In *R: A Language and Environment for Statistical Computing*, Version 3010.82: Revision 5504. Available via http://www.rmetrics.org [accessed July 14, 2014.

AUTHOR BIOGRAPHIES

JAMES R. THOMPSON currently works at MITRE Corporation in Washington, D.C. He is a member of INFORMS, and his research interests are focussed on stochastic modeling techniques, including agent-based simulation, large-scale discrete-event simulation, and system dynamics as applied to problems of policy making and decision analysis in the context of government, economics, and healthcare. His email address is jrthompson@mitre.org.

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, and ASEE; 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.