Proceedings of the 2019 Winter Simulation Conference
N. Mustafee, K.-H.G. Bae, S. Lazarova-Molnar, M. Rabe, C. Szabo, P. Haas, and Y.-J. Son, eds.

SIMULATING MULTIFRACTAL SIGNALS FOR RISK ASSESSMENT

Damon Frezza
James Thompson
David Slater
Garry Jacyna

The MITRE Corporation
7525 Colshire Dr
McLean, VA 22102, UNITED STATES

ABSTRACT

Many data sets collected from physical processes or human engineered systems exhibit self-similar properties that are best understood from the perspective of multifractals. These signals fail to satisfy the mathematical definition of stationarity and are therefore incompatible with Gaussian-based analysis. Efficient algorithms for analyzing the multifractal properties exist, but there is a need to simulate signals that exhibit the same multifractal spectrum as an empirical data set. The following work outlines two different algorithms for simulating multifractal signals and addresses the strengths and weaknesses of each approach. We introduce a procedure for fitting the parameters of a multifractal spectrum to one extracted empirically from data and illustrate how the algorithms can be employed to simulate potential future paths of a multifractal process. We illustrate the procedure using a high-frequency sample of IBM's stock price and demonstrate the utility of simulating multifractals in risk management.

1 INTRODUCTION

A large number of empirical data sets collected as signals from natural phenomena have been shown to exhibit self-similar properties that are best understood from the perspective of fractals and multifractals. Examples include electrocardiogram time series from the human heart (Shekatkar et al. 2017), financial time series (Thompson and Wilson 2016), Internet traffic (Chakraborty et al. 2004), turbulence (Frisch and Parisi 1980), and the sequence of nucleotides that constitute our DNA (Moreno et al. 2011). These signals fail to satisfy the mathematical definition of stationarity and are therefore incompatible with Gaussian-based analysis (Shumway and Stoffer 2006). We must therefore turn our attention to the power-laws that govern the statistics of self-similar processes.

1.1 Multifractals

A stochastic process $X(t) \in \mathbb{R}$ is said to be self-similar with Hurst exponent $H \in [0,1)$ if for any $c > 0$ and support t, $E[|X(t+s) - X(t)|^q] \stackrel{d}{=} c(q)s^{qH}$ where the d over the equals sign indicates equal in distribution and $q \in \mathbb{R}$ is the moment parameter (Mandelbrot and van Ness 1968). Mandelbrot and Van Ness showed that the Hurst exponent H in the above definition plays a key role in the autocorrelation function of self-similar processes. Specifically when $0 < H < 0.5$ the process exhibits short-range dependence and is anti-persistent. Conversely when $0.5 < H < 1.0$, the process exhibits long-range dependence and is persistent. Finally, when $H = 0.5$, the increments of the stochastic process have zero autocorrelation.

The concept of multifractals generalizes the above definition such that increments of different sizes are governed by different power-laws. Thus, instead of a global scalar exponent H there is a spectrum of local Hölder exponents $h_q(t)$. In order to introduce multifractals, we must first define Hölder regularity.

Formally, a process is *Hölder regular* if in the limit it approaches a differentiable polynomial. In this case, we can define the Hölder exponent $h_q(t)$ as a measure of regularity if there is a polynomial P of degree less than α such that $|X(t)-P(t-t_0)| \leq c|t-t_0|^{\alpha}$ where c is a constant greater than 0. Then the Hölder exponent is the supremum α-value for which $h_q(t) = \sup\{\alpha \geq 0 : |X(t)-P(t-t_0)| \leq c|t-t_0|^{\alpha} \text{ as } |t-t_0| \to 0\}$. As such, the Hölder exponent is a measure of local regularity about the point t (Ouahabi 2012). Given this definition of local regularity, we say that a process is *multifractal* if $E[|X(t+s)-X(t)|^q] = c(q)s^{\tau(q)+1}$. The function $\tau(q)$ is often referred to as the scaling function and it replaces the scalar H in our definition of self-similar processes. The collection of points that exhibit the same Hölder exponent h_q are not sub-intervals of a multifractal process, but rather they are scattered throughout the support. If we define this set of iso-Hölder points as $\mathscr{T} = \{t : h_q(t) = h_q\}$ then it can be shown that the set $\mathscr{T}$ constitutes a fractal set with fractal dimension D_h (Mandelbrot 1989). Furthermore, the definition of a multifractal measure ties the fractal dimension of an iso-Hölder set to the scaling function $\tau(q)$ via the Legendre transform (Harte 2001). That is, $D_h = \min_q\{qh_q - \tau(q)\}$. Solving for D_h, it is easy to see that $h_q = \frac{d}{dq}\tau(q)|_{q=q_o}$ where q_o is the value of q that produces the minimum. Thus, the fractal dimension D_h is a relative measure of the prevalence of a given Hölder expontent h_q in a multifractal process with scaling function $\tau(q)$.

The objective of multifractal analysis is to quantify, in comparable finite terms, the properties of data sets that exhibit this behavior. In this way we can compare complex systems to one another and thus measure the relative complexity of given systems in various states. Extensive research of different algorithms for analyzing the multifractal spectrum has led to efficient implementations that can quickly and accurately extract an empirical spectrum from data (Jaffard, Lashermes, and Abry 2006). However, there is a need to be able to simulate signals that exhibit the same multifractal spectrum as the empirical data set.

1.2 Cascades

One way to simulate multifractal signals is to use a multiplicative cascade process. As the name implies, the procedure involves a cascade from coarse scales to finer scales through a set of multipliers at each scale. One controls the properties of the signal by varying either the method of cascade from one scale to the next or the distribution of the multipliers at each scale or both. This level of control allows for a wide variety of signals to be created. The resulting multiplicative cascades are multifractal with positive increments. Early cascade methods study extensively by Mandelbrot restricted the method of cascade to integer-based branching at each subsequent scale (Mandelbrot 1974). These Canonical Mandelbrot Cascades have three major drawbacks:

1. The iterative construction results in discrete scale invariance.
2. The increments are not stationary.
3. Strictly positive increments lead to a monotonically increasing signal, limiting the diversity of signals that can be simulated.

Several approaches have been proposed to overcome these limitations. The method we focus on is the Compound Poisson Cascade (CPC). The Poisson point process that underlies the construction of CPCs removes the reliance on deterministic geometric grids used in previous constructions. This addresses the mechanism responsible for discrete scale invariance and non-stationary increments. CPC increments are still strictly positive so, to address the third drawback and thereby obtain a wider selection of simulated signals, we embed the multifractal cascade into fractional Brownian motion (fBm) as the *time* or *support* parameter (Calvet and Fisher 2002). The result is a well-behaved multifractal signal with positive and negative increments. This process can then be used to simulate observed signals and perform risk management analysis.

1.3 Wavelet Leaders

In this section we give a brief overview of the algorithms used to extract multifractal spectra from empirical data sets. More comprehensive descriptions can be found in Jaffard, Lashermes, and Abry (2006) and Wendt (2008). Wavelet-based analysis relates the coefficients from wavelet decomposition to the scaling function $\tau(q)$. A function $\psi(t)$ is a wavelet function if it exhibits a short oscillation with an amplitude that starts at zero and returns to zero with finite support. If the wavelet function $\psi(t)$ is in $\mathbb{L}^2(\mathbb{R})$ and the functions $\psi_{j,k}(t) = 2^{j/2}\psi(2^j t - k)$ for $j \in \mathbb{Z}$ and $k \in \mathbb{Z}$ create an $\mathbb{L}^2(\mathbb{R})$ orthonormal basis then this system of functions is said to be an orthonormal wavelet basis. Functions $f(x)$ that are in the $\mathbb{L}^2(\mathbb{R})$ space can be written as a decomposition of the orthonormal basis functions $f(x) = \sum_{j,k} c_x(2^j,k)\psi(2^j x - k)$ where the wavelet coefficients are defined as $c_x(2^j,k) = 2^j \int f(x)\psi(2^j x - k)dx$.

Note that if the wavelet function $\psi(t)$ is m-times continuously differentiable, then its first m moments integrate to zero. Due to the vanishing moments, if $X(t)$ exhibits polynomial behavior of a degree less than or equal to m, the wavelet coefficients are detrended transformations and do not carry that polynomial information (Ouahabi 2012). Given a wavelet decomposition at many scales 2^j we can define the statistical moments of order q using the structure function $S(2^j,q) = \frac{1}{n_j}\sum_{k=1}^{n_j} |c_x(2^j,k)|^q$ where n_j is the quantity of translations k at scale 2^j, $q \in \mathscr{Q}$, and $[0,1] \subseteq \mathscr{Q}$. In the context of the multifractal formalism, wavelet-based analysis then asks whether the structure function scales according to a power-law. That is, if the process is multifractal then $S(2^j,q) \sim 2^{j\tau(q)}$ as $2^j \to 0$, where $\tau(q)$ is the scaling function from the definition of multifractals. Given $\tau(q)$ we need only to apply the Legendre transform to obtain the multifractal spectrum.

This work introduces a new procedure for fitting the parameters of a multiplicative cascade process to the multifractal properties extracted from empirical data. We show how this procedure can be employed for risk management and we analyze its performance on a financial time series to illustrate its utility. Section 2 describes the cascade algorithms and outlines our procedure for fitting and simulating multifractals that match observed signals. In section 3 we describe the data we used for the experiment and outline our experimental analysis. We also use these results in a forecasting exercise to illustrate the application to risk management before concluding in section 4.

2 METHODS

2.1 Compound Poisson Cascades

As mentioned in section 1.2, Canonical Mandelbrot cascades have several limitations. We therefore restrict our attention in this work to Compound Poisson Cascades (CPCs). CPCs require strictly positive multipliers to create a multifractal process. We therefore examine log-Normal multipliers of the form $\mathscr{W} = e^X$, where $X \sim \mathscr{N}\left(\mu, \sigma^2\right)$. Following the description outlined by Wendt (2008), the CPC signal is defined as:

$$Q_r(t) = c \prod_{(t_i,r_i)\in C_r(t)} W_i \tag{1}$$

where (t_i, r_i) are the positions of a 2-dimensional Poisson point process with intensity measure $dm(t,r)$ supported on the rectangle $I = \{(t',r') : r_{min} \le r' \le 1, -\frac{1}{2} \le t' \le T + \frac{1}{2}\}$. The W_i are the associated random variables sampled from $\mathscr{W}$. We define the cone of influence for point t as $C_r(t) = \{(t',r') : 0 \le r' \le 1, t - \frac{r'}{2} \le t' \le t + \frac{r'}{2}\}$. The normalizing constant c is defined such that $\mathbb{E}[Q_r(t)] = 1$. Algorithm 1 outlines the simulation procedure for CPCs. Note that we present one option for the intensity measure $dm(t,r) = \frac{drdt}{r^2}$ that guarantees stationarity. The intensity measure is accounted for when calculating $\vec{t}$ and $\vec{r}$ in Algorithm 1. The construction procedure and resulting signals are also illustrated in Figure 1 that was adapted from Wendt (2008).

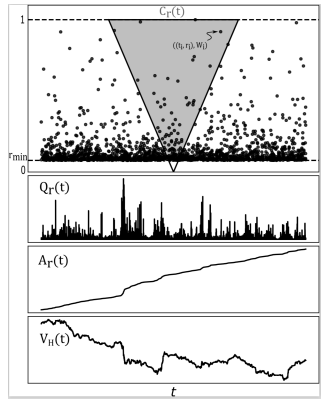

Figure 1: The cone $C_r(t)$ and multipliers $W_{i,j}$, the signal $Q_r(t)$, $A_r(t)$ the integral of $Q_r(t)$, and $V_H(t)$ that results from combining $A_r(t)$ with fractional Brownian motion with Hurst parameter H. Based on a figure in Wendt (2008).

Algorithm 1: Constructs a CPC with multipliers from a given distribution

Input:
n = length of signal (power of 2)
$\mathscr{W}$ = the distrubtion for the random multipliers
r_{min} = the minimum resolution for the signal
T = max time value
Output:
$\vec{Q}$ = the CPC signal of length n

$\lambda \leftarrow (T+1)\left(\frac{1}{r_{min}} - 1\right)$
$N \leftarrow$ 1 realization from $Poisson(\lambda)$
$u_{max} \leftarrow \frac{1}{r_{min}} - 1$
$\vec{u}_{init} \leftarrow \vec{0} : \|\vec{u}_{init}\| = T+1$
$\vec{u}_{rand} \leftarrow$ vector of N samples from $Uniform(0, u_{max})$
$\vec{u} \leftarrow [\vec{u}_{init}\ \vec{u}_{rand}]$
$\vec{r} \leftarrow \frac{1}{1+\vec{u}}$
$\vec{t} \leftarrow$ vector of $N+T+1$ samples from $Uniform(-\frac{1}{2}, T-\frac{1}{2})$
$\vec{W} \leftarrow$ vector of $N+T+1$ samples from $\mathscr{W}$
$\vec{Q} \leftarrow \vec{1} : \|\vec{Q}\| = n$
$\vec{x} \leftarrow (0, \frac{T}{n-1}, \frac{2T}{n-1}, \frac{3T}{n-1}, ..., T)$
for $i \leftarrow 1$ ***to*** $N+T+1$ **do**
  $\vec{J} \leftarrow which\left(|\vec{x} - t_i| \leq \frac{r_i}{2}\right)$
  foreach j ***in*** $\vec{J}$ **do**
    $Q_j \leftarrow W_j Q_j$
$c \leftarrow r_{min}^{exp(\mu + \frac{\sigma^2}{2}) - 1}$
$\vec{Q} \leftarrow c\vec{Q}$
return $\vec{Q}$

Algorithm 2: Combines a fractional Brownian motion signal with CPM time

Input:
$\vec{A}_r$ = a CPM signal of length n, used for time increments
H = the Hurst parameter for the fractional Brownian motion
d_t = the time sampling rate
o_t = the over sampling rate
Output:
$\vec{V}_H$ = the fractional Brownian motion signal embedded in CPM time

$z \leftarrow max(\vec{A}_r) d_t$
$N_{fBm} \leftarrow 1 + 2^{ceiling(\log_2 z)}$
$\vec{B}_H \leftarrow$ an fBm signal of length N_{fBm} with Hurst param H
$\vec{B}_H \leftarrow z^H \vec{B}_H$
$\vec{i} \leftarrow round(\vec{A}_r o_t)$
$i_1 \leftarrow max(i_1, 1)$
$\vec{V}_H \leftarrow \left(b_i \in \vec{B}_H | i \in \vec{i}\right)$
return $\vec{V}_H$

In Equation (1) we introduced a normalizing constant c such that $\mathbb{E}[\vec{Q}] = 1$. For a CPC with log-Normal multipliers (CPC-LN), $c = r_{min}^{exp(\mu + \frac{\sigma^2}{2}) - 1}$ where the final resolution $r_{min} = \frac{1}{n}$. Using these parameters, we

can characterize the multifractal spectrum and its log-cumulants. Equation (2) is the scaling function for a CPC-LN that — through the Legendre transform — defines the multifractal spectrum in Equation (4). Because Equation (4) has no closed form solution, we define the expression $\Lambda(h,q)$ and take the derivative with respect to q in Equation (5). Setting $\frac{\partial\Lambda}{\partial q}=0$ and solving numerically for q at a given value of h yields $D(h)$. Since the multifractal spectrum can be fully described by its log-cumulants (see Wendt (2008)), we provide (6) that is used to calculate c_p for all p.

$$\tau(q) = (\eta-1)q + 1 - e^{q\mu+\frac{q^2\sigma^2}{2}} + qe^{\mu+\frac{\sigma^2}{2}} \tag{2}$$

$$D(h) = \min_{q\neq 0}(d + qh - \tau(q)) \tag{3}$$

$$D(h) = \min_{q\neq 0}\underbrace{\left(d + q(h+1-\eta) - 1 + e^{q\mu+\frac{q^2\sigma^2}{2}} - qe^{\mu+\frac{\sigma^2}{2}}\right)}_{\Lambda(h,q)} \tag{4}$$

$$\frac{\partial\Lambda(h,q)}{\partial q} = h + 1 - \eta + e^{q\mu+\frac{q^2\sigma^2}{2}}\left(\mu + q\sigma^2\right) - e^{\mu+\frac{\sigma^2}{2}} \Rightarrow 0 \tag{5}$$

$$c_p = \frac{d^p}{dq^p}\tau(q)|_{q=0} \tag{6}$$

2.2 Fractional Brownian Motion in CPM Time

In this section we address the issue of strictly positive increments in the multifractal signal generated from CPCs. The diversity of signals that can be simulated with only positive increments is limited, but by allowing the support of a fractional Brownian motion (fBm) process to be represented by a multiplicative cascade we can simulate a wider variety of signals (Calvet and Fisher 2002). We start by formally defining Compound Poisson Motion (CPM) as $A(t) = \lim_{r\to 0}\int_0^t Q_r(s)ds$. The approximation $A_r(t)$ to the limit integral that defines CPM is just the cumulative sum of $Q_r(t)$. To construct a multifractal process with positive and negative increments $V_H(t)$, we combine an fBm signal $B_H(t)$ with CPM *time*. The resultant multifractal is well-defined and inherits properties from both the original fBm and CPM signals. The scaling behavior of $V_H(t)$ is a direct consequence of the self-similarity of fBm combined with the scaling behavior of CPM (Chainais et al. 2005; Pipiras and Taqqu 2017). The resulting signal has the rich interwoven scaling characteristics of the CPM, as well as the added self-similarity of the fBm. Algorithm 2 outlines the procedure for constructing the combined process and the bottom right panel of Figure 1 illustrates a sample process path of $V_H(t)$.

As with the multifractal spectrum for the CPC, $D(h)$ for this process has no closed-form solution. We therefore derive the equation for the scaling function $\tau(q)$ that facilitates the numerical approximation of $D(h)$ via the Legendre transform. Using these equations and the procedure we introduce in the rest of this section we can simulate replications of $V_H(t)$ with optimal parameters μ^*, σ^{*2}, η^*, and H^*.

$$\tau(q) = q\eta + 1 - e^{qH\mu+\frac{q^2H^2\sigma^2}{2}} + qHe^{\mu+\frac{\sigma^2}{2}} \tag{7}$$

$$D(h) = \min_{q\neq 0}(d + qh - \tau(q)) \tag{8}$$

$$D(h) = \min_{q\neq 0}\underbrace{\left(d + q(h-\eta) - 1 + e^{qH\mu+\frac{q^2H^2\sigma^2}{2}} - qHe^{\mu+\frac{\sigma^2}{2}}\right)}_{\Lambda(h,q)} \tag{9}$$

$$\frac{\partial\Lambda(h,q)}{\partial q} = h - \eta + e^{qH\mu+\frac{q^2H^2\sigma^2}{2}}\left(\mu H + q\sigma^2H^2\right) - He^{\mu+\frac{\sigma^2}{2}} \Rightarrow 0 \tag{10}$$

$$c_p = \frac{d^p}{dq^p}\tau(q)|_{q=0} \tag{11}$$

Equation (7) — derived directly from Equation (2) and the self-similarity of $B_H(t)$ — shows how the $V_H(t)$ process scales with support t. Note that although one can calculate the log-cumulants directly using Equation (11) with the relevant scaling function defined in Equation (7), it can also be shown that $c_p^{V_H} = H^p c_p^{A_r}$ for all $p > 1$, where $c_p^{A_r}$ are the log-cumulants of the CPM process used for the time increments and H is the Hurst exponent of the $B_H(t)$ process. Note that c_1 has a similar relationship with H, but also contains the fractional integration parameter η.

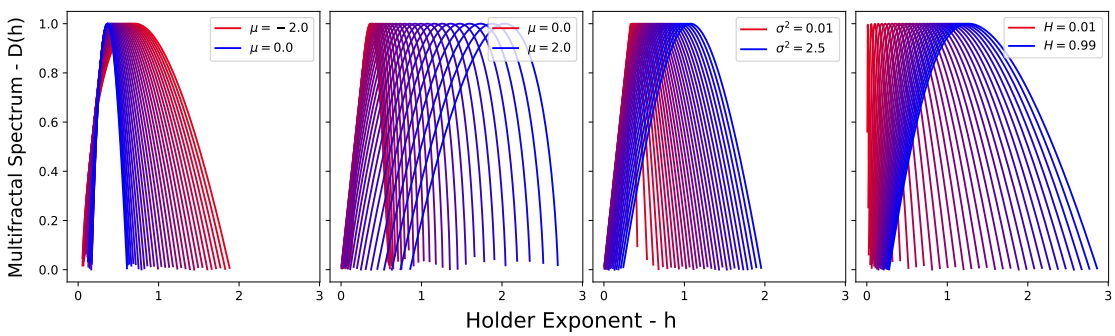

Figure 2: The different multifractal spectra of $V_H(t)$ achieved by varying one parameter and fixing the other two. Each plot has 30 curves based on an evenly spaced parameter sweep between the lower and upper values shown.

2.3 Fitting the Multifractal Properties of an Empirical Signal

Given the flexibility of CPCs combined with fBm, they are well-suited for fitting a wide variety of empirical signals. Figure 2 illustrates how the μ, σ^2, and H parameters affect the shape of the resultant multifractal spectrum and gives a visualization of the range of signals that can be simulated. In addition we can vary the fractional integration parameter η to shift the curve left for $\eta < 0$ and right for $\eta > 0$ without impacting the shape. We use these four parameters to fit $V_H(t)$ to empirical data.

The shape of an empirically extracted multifractal spectrum is often difficult to match exactly. To find the best fit we compare the use of two options: $\tau(q)$ from Equation (7) and $D(h)$ from Equations (9) and (10). In each case, we compare the theoretical values for a simulated signal $V_H(t)$ to the empirical values extracted from data using the wavelet leaders approach described in section 1.3. There are approaches to fitting multifractal properties that focus on matching the log-cumulants of the scaling function. In complex cases, the curve fitting approaches for $\tau(q)$ and $D(h)$ tend to fail more frequently. Algorithms for these sorts of approaches have been explored by other researchers and are beyond the scope of the current work (Wendt, Abry, and Jaffard 2007).

2.3.1 Fitting the Scaling Function $\tau(q)$

The scaling function $\tau(q)$ is often easier to fit because it can have long sections that are close to linear. Additionally, $\tau(q)$ has a closed-form solution for $V_H(t)$, which simplifies the fitting procedure and makes it accurate and efficient. For this procedure, we utilize the Broyden-Fletcher-Goldfarb-Shanno (BFGS) Hessian Update Strategy. We start by defining the objective function in terms of the closed-form solution for $\tau(q)$:

$$Z_\tau = \frac{1}{n_q}\sum_{q=q_1}^{q_n} w_q\left(\tau(q) - \hat{\tau}(q)\right)^2 \tag{12}$$

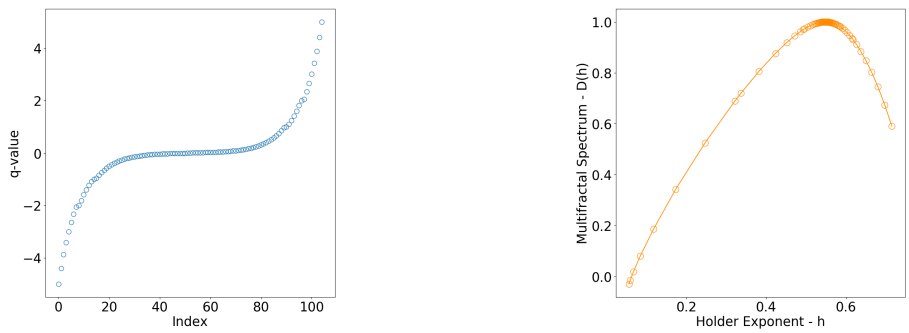

(a) The logarithmically spaced q-values (b) The Hölder exponents h created from the log-spaced q-values

Figure 3: Moment orders q in the wavelet leaders algorithm.

where $\tau(q)$ is Equation (7), $\hat{\tau}(q)$ is the estimated scaling function of the data, $\{q_1, \ldots, q_n\}$ are the log-spaced moment orders used to calculate $\hat{\tau}(q)$, w_q is the weight for the qth moment order, and n_q is the total number of moment order values. We include w_q for more fine tuning in the fitting procedure. We begin with $w_q = 1$ for all q, which weights all points equally. If a portion of the fit deviates substantially from the empirical estimate $\hat{\tau}(q)$, we weight that section of the curve higher to increase the corresponding portion of Equation (12). The ultimate impact is to emphasize deviations in that section of the curve. This objective function is then minimized using BFGS to find the optimal parameters μ^*, σ^{*2}, H^*, η^* which define the optimal fitted process $V_H^*(t)$. BFGS allows for bounded optimization which is necessary to ensure the following conditions:

$$\sigma^2 > 0 \tag{13}$$

$$0 < H < 1 \tag{14}$$

$$\mu + \sigma^2 < \log\sqrt{2} \tag{15}$$

Note we incorporate the condition in Equation (15) because CPCs have a level of degeneracy for certain choices of μ and σ^2. Equation (15) ensures a stable process and effective simulation.

2.3.2 Fitting the Multifractal Spectrum $D(h)$

Unlike the scaling function, the multifractal spectrum $D(h)$ has no closed-form for the model CPC process $V_H(t)$. As a result, fitting $D(h)$ from Equation (9) requires a more complex procedure than fitting $\tau(q)$. To use BFGS to optimally fit $D(h)$ to the spectrum of the data $\hat{D}(h)$ we must now define the objective function in terms of the numerical approximation of the Legendre transform:

$$Z_D = \frac{1}{n_h} \sum_{h=h_1}^{h_n} w_h \left(\tilde{D}(h) - \hat{D}(h)\right)^2 \tag{16}$$

where $\tilde{D}(h)$ is the numerical approximation of $D(h)$, $\hat{D}(h)$ is the estimate of the empirical multifractal spectrum, $\{h_1, \ldots, h_n\}$ are the Hölder exponents used to calculate $\hat{D}(h)$, w_h is the weight for the given Hölder exponent h, and n_h is the total number of Hölder exponents.

To numerically approximate $D(h)$ we must first find the q-value that minimizes $\Lambda(h,q)$ as defined in Equation (9). Recall that the q-values are predetermined as input parameters to the wavelet leaders algorithm (see Figure 3a). However, better fits can be achieved if the minimum q-values are extracted from the resulting multifractal spectrum $\hat{D}(h)$. This amounts to solving for q-values that fall along the extrapolated curve shown in Figure 3b. We thus set $\frac{\partial \Lambda(h,q)}{\partial q} = 0$ and solve for q as shown in Equation (10). We use the Brent root-finding algorithm to find the minimizing q^* for each h (Brent 1971). Because

$\Lambda(h,q)$ depends on the decision parameters μ, σ^2, H, η, we must re-solve for the q^* values during each evaluation of the BFGS algorithm as detailed in Algorithm 3. As with fitting $\tau(q)$, the bounds in (13) - (15) are enforced to ensure valid solutions.

Algorithm 3: Fitting $V_H(t)$ to empirical data using $D(h)$.

Input:
$\vec{h}$ = the vector of Hölder exponents associated with the $\hat{D}(h)$ estimates from the empirical data.
$\hat{D}(h)$ = the multifractal spectrum estimates for each $h \in \vec{h}$.
Output:
Θ^* = the optimal parameters $\left(\mu^*, \sigma^{*2}, H^*, \eta^*\right)$ of the best fit $V_H^*(t)$ process.

Define: $\Theta \leftarrow \left(\mu, \sigma^2, H, \eta\right)$
Define: $\Lambda_h \leftarrow \frac{\partial \Lambda(h,q)}{\partial q}$ for a given h
Initialize Θ, Z_D
while *BFGS suboptimal* **do**
  $\Theta \leftarrow$ BFGS_Parameter_Update(Z_D, Θ)
  $\tilde{D}(h) \leftarrow \{\}$
  foreach $h \in \vec{h}$ **do**
    $q_h^* \leftarrow$ Brent_Root(Λ_h, Θ)
    $\tilde{D}(h) \leftarrow \{\tilde{D}(h), \Lambda(h, q_h^*)\}$
  $Z_D \leftarrow \frac{1}{n_h} \sum w_h \left(\tilde{D}(h) - \hat{D}(h)\right)^2$
return Θ^*

Algorithm 4: Scaling $v(t)$, the simulated increments of the $V_H(t)$ process.

Input:
$X(t)$ = the empirical data.
$v(t)$ = the signals simulated from the $V_H(t)$ process where $v_i(t)$ is the ith replication.
Output:
$\tilde{v}(t)$ = the scaled simulated signals.

$X'(t) \leftarrow X(\vec{t}) - X(\vec{t} - \vec{1})$ # first difference of $X(t)$
$S_v \leftarrow \{\}$
foreach $i = 1 : n_c$ **do**
  $v_i'(t) \leftarrow v_i(\vec{t}) - v_i(\vec{t} - \vec{1})$ # first difference of $v_i(t)$
  $S_v \leftarrow \{S_v, v_i'(t)\}$
foreach $i = 1 : n_c$ **do**
  $\tilde{v}_i(t) \leftarrow CumSum\left(\frac{v_i'(t)}{\langle |S_v| \rangle} \langle |X'(t)| \rangle\right)$
return $\tilde{v}(t)$

2.3.3 Deciding between Methods

The decision between the two methods depends on the intended application of the results. Figure 4 shows the multifractal spectra resulting from fitting $D(h)$ (top left) and fitting $\tau(q)$ (bottom left). Note that most of the deviation in the $D(h)$ fit occurs towards the center of the spectrum while most of the deviation from the $\tau(q)$ fit is in the tails of the spectrum. The peak and center of the spectrum are representative of the average behavior of the underlying process. The deviations in the tails of spectrum are the result of rare events. The careful practitioner should weigh these characteristics when deciding between the two methods. If average behavior is more important to their current analysis, fitting $\tau(q)$ is the more prudent method. If, on the other hand, extreme events are the greater concern, fitting $D(h)$ is advisable.

2.4 Scaling Simulated Increments

The multifractal spectrum of a signal is independent of the magnitude of its increments such that $X(t)$ and $cX(t)$ have the same multifractal spectrum if $X(t)$ is a multifractal and c is a scalar. As a result, another aspect of fitting a cascade to data is the appropriate scaling of increments. After finding the optimal parameters that define the best fit $V_H^*(t)$ process, we simulate n_c cascades. The increments of these cascades are all on the same scale as one another but not the same scale as the increments of the data. The goal of the simulation of these cascades is improving risk management and the model's ability to encapsulate rare events. Therefore, to preserve the rare events and simultaneously scale the increment magnitudes to an appropriate level, we scale by the average magnitude of all of the cascade increments simulated and the average magnitude of the increments of the empirical data as shown in Algorithm 4.

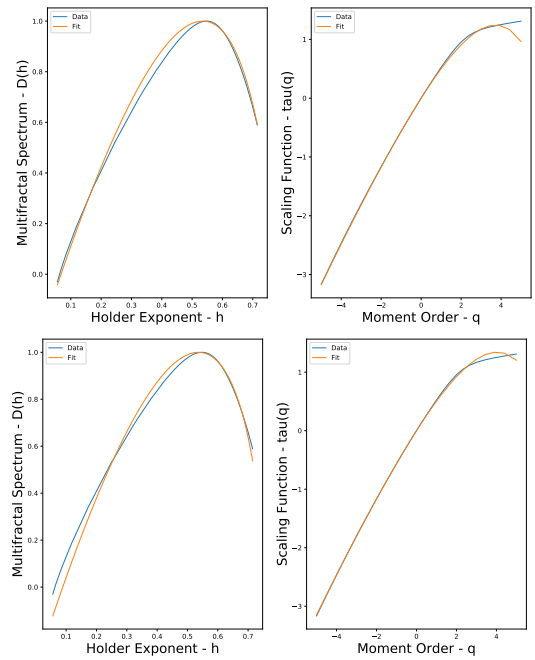

Figure 4: The multifractal spectrum (left) and scaling function (right) from an empirical data set (blue) and a simulated optimal fit $V_H^*(t)$ (orange) obtained by fitting $\hat{D}(h)$ (top) or $\hat{\tau}(q)$ (bottom).

3 EXPERIMENTS

3.1 Data

In this section we analyze a high-frequency data set from the New York Stock Exchange taken from the Wharton Research Data Services (WRDS) data (Wharton 2013). WRDS has compiled trade and quote (TAQ) data from various exchanges into a database to facilitate fine-scale analysis. We arbitrarily selected IBM for the time frame covered by our limited data set; January 2, 1998 to December 31, 2003. The TAQ data records every trade and quote submitted to the exchange for a given asset. The trades constitute a change in price and when concatenated together represent a higher frequency price time series than just the daily closing prices. During the time frame covered by our data, online trading was just getting started so the average daily trades steadily increased from 1,310 trades per day in 1998 to over 4,445 trades per day by the end of 2003. Our data set consists of 4,393,502 timestamped trades.

3.2 Fitting and Simulating

The objective of our experiment was to define a process $V_H^*(t)$ with multifractal properties that most closely matched the empirical scaling function $\hat{\tau}(q)$ extracted from the IBM log price using the wavelet leaders algorithm. The experiment employed the procedure described in section 2.3 to find the optimal parameters μ^*, σ^{*2}, H^*, and η^*. The resulting fit is shown across the bottom of Figure 4 and has optimal parameters $\mu^* = 0.296$, $\sigma^{*2} = 0.05$, $H^* = 0.654$, and $\eta^* = -0.158$.

The scaling function $\hat{\tau}(q)$ (and by extension the multifractal spectrum $\hat{D}(h)$) completely describes how the moments of the stochastic price process change across scales. In order to put the optimal fit to practical use it is convenient to simulate many instantiations of the optimal process $V_H^*(t)$. Using the algorithms from section 2 and the optimal parameters for $V_H^*(t)$, we generated 30 independent CPCs and combined them with five fBm signals for a total of 150 replications of $V_H^*(t)$. Matching the multifractal spectrum ensures that the relative increments of the simulated signals match the original multifractal properties of the empirical data. However, it does not ensure that the magnitude of the simulated increments is representative

of the empirical signal. We therefore scaled each simulated $V_H^*(t)$ according to the procedure outlined in section 2.4 such that the average increment magnitude of all 150 signals was equal to the average increment magnitude of the IBM signal as described. Figure 5 shows the data increments (left) and the corresponding scaled cascade increments (right).

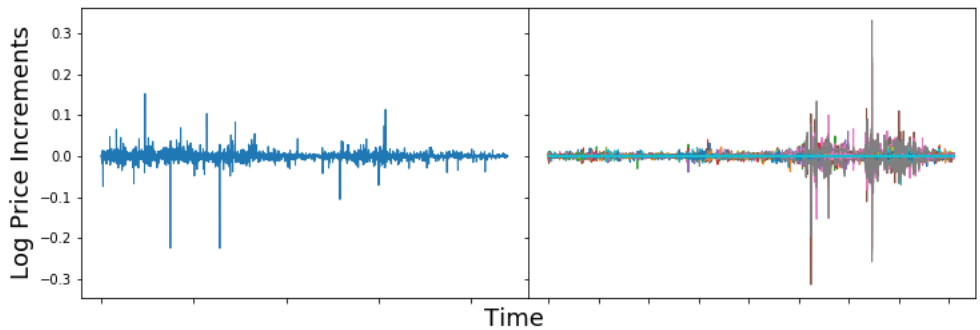

Figure 5: Left: the increments of the log price of IBM. Right: the scaled increments of the 150 simulated signals $V_H(t)$.

For illustrative purposes we use the simulated signals to create a projection cloud and a 95% prediction interval to forecast the future log price of IBM as shown in Figure 6. An important feature of this prediction interval is the exponential expansion of the bounds. Traditional statistical models often do not account for long range rare events that can occur more frequently when the moments exhibit power-law relationships with scale. The right-hand side of Figure 6 shows an overlay of the actual daily closing price of IBM (black) during the two months that follow the last trade in our data set. It is clear that the 95% prediction interval encapsulates the movements in the closing price over the period considered and captures the increasing risk of larger changes in price over longer forecasts.

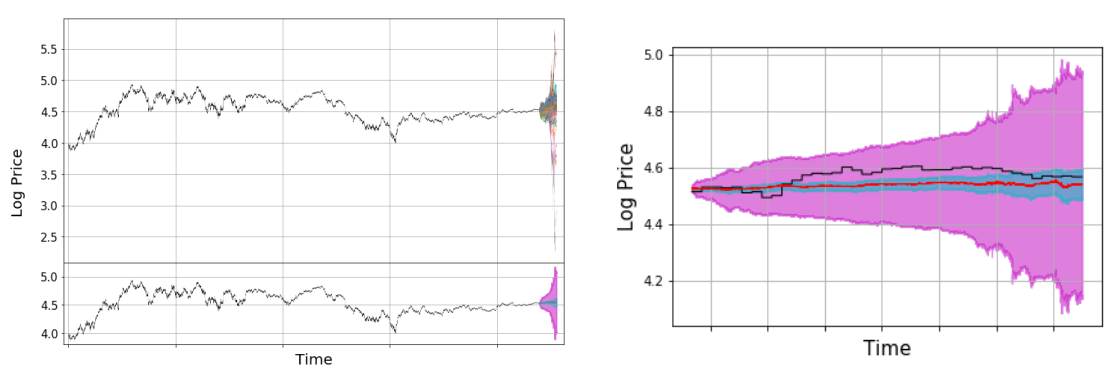

Figure 6: IBM log price with: (Top left) simulated $V_H(t)$ and (Bottom left and right) the 95% confidence interval (cyan) and 95% prediction interval (magenta).

3.3 Simulation Burn-in and Fractional Integration

As with most limit processes, fractional integration requires a sufficient number of data points to compute the integrated values around a given support point. It follows that executing a simulation procedure that requires fractional integration will require some degree of *burn-in* before the accuracy of the results become reliable. We therefore incorporate a burn-in period for each simulation to allow the signal to converge.

Figure 7 shows how the initial values are unusable and gives an indication of what value of burn-in might be prudent. To further investigate the amount of burn-in required, we examine the box plots in Figure 8 that compare the estimated log-cumulants across different burn-in periods. Comparison with the "no burn-in" case illustrates the bias inherent in estimates that do not employ a burn-in period.

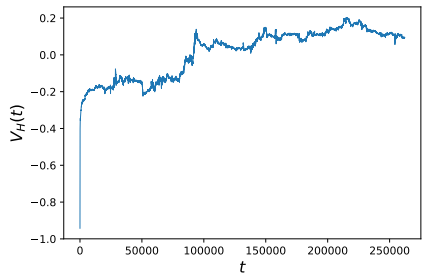

Figure 7: A single simulation of $V_H^*(t)$ illustrating the necessity of a burn-in period.

Figure 8 also shows the accuracy of the $V_H^*(t)$ simulation compared with the empirically extracted log-cumulants. For a 50,000 point burn-in, the distribution of c_1 is centered on the empirically estimated value. The distribution for c_2 is negatively biased from the target value indicating that we should expect the simulated signals to have slightly wider multifractal spectra than that of the data. Wider spectra exhibit more power-laws and thus constitute a more conservative estimate in terms of forecasting.

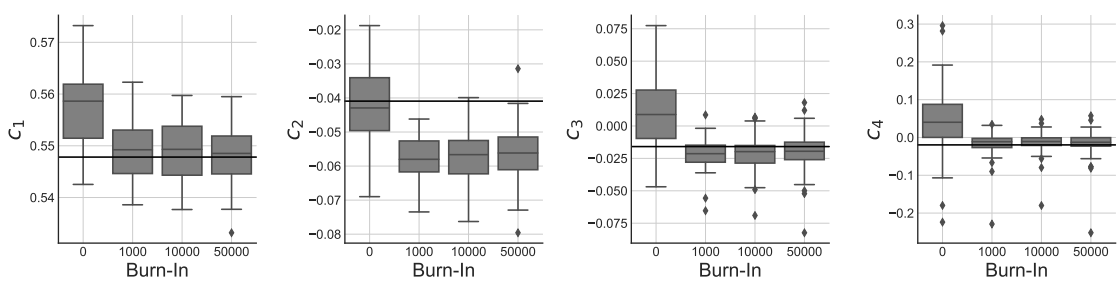

Figure 8: The log-cumulants of the IBM log price data (black) compared with $V_H^*(t)$ at different values of burn-in.

4 CONCLUSION

Advances in computational science have led to an explosion in the collection, archiving, and analysis of data. Many of these data sets exhibit heavy-tails in the distribution of increments, auto-correlation functions with seemingly infinite lag order, and bursts of volatility that belie the assumptions of stationarity. These data sets are dominated by power-laws and self-similarity that are best analyzed through the lens of multifractals. We now have efficient and accurate algorithms for extracting the multifractal properties from data, but in order for that information to be of practical use we need to be employ those properties in risk management tools such as forecasting and threshold modeling.

In this work we introduced a procedure for statistically matching the multifractal properties of an empirical data set with a robust model for simulating multifractals. We illustrated the utility of the procedure by fitting the parameters of IBM's trade and quote time series from 1998 to 2001 and using the simulations resulting from the procedure to forecast the future price path of the asset. This demonstrated how this technique can be used to create prediction intervals around notoriously non-stationary time series and thus benefit risk managers.

ACKNOWLEDGMENTS

Approved for Public Release; Distribution Unlimited. Public Release Case Number 19-2036.

REFERENCES

Brent, R. P. 1971. "An Algorithm with Guaranteed Convergence for Finding a Zero of a Function". *The Computer Journal* 14(4):422–425.

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

Chainais, P., R. Riedi, and P. Abry. 2005. "Warped Infinitely Divisible Cascades: Beyond Power Laws". *Traitement du Signal* 22(1):27–39.

Chakraborty, D., A. Ashir, T. Suganuma, G. Mansfield Keeni, T. K. Roy, and N. Shiratori. 2004. "Self-Similar and Fractal Nature of Internet Traffic". *International Journal of Network Management* 14(2):119–129.

Frisch, U., and G. Parisi. 1980. "Fully Developed Turbulence and Intermittency". *New York Academy of Sciences, Annals* 357:359–367.

Harte, D. 2001. *Multifractals: Theory and Applications*. Boca Raton, FL: Chapman and Hall/CRC.

Jaffard, S., B. Lashermes, and P. Abry. 2006. "Wavelet Leaders in Multifractal Analysis". In *Wavelet analysis and applications*, edited by M. I. V. Tao Qian and X. Yuesheng, 201–246. Switzerland: Birkhäuser, Basel.

Mandelbrot, B. B. 1974. "Intermittent Turbulence in Self-similar Cascades: Divergence of High Moments and Dimension of the Carrier". *Journal of Fluid Mechanics* 62(2):331–358.

Mandelbrot, B. B. 1989. "Multifractal Measures, Especially for the Geophysicist". In *Fractals in geophysics*, 5–42. Switzerland: Birkhäuser, Basel.

Mandelbrot, B. B., and J. W. van Ness. 1968. "Fractional Brownian Motions, Fractional Noises and Applications". *Society for Industrial and Applied Mathematics Review* 10(4):422–437.

Moreno, P. A., P. E. Vélez, E. Martínez, L. E. Garreta, N. Díaz, S. Amador, I. Tischer, J. M. Gutiérrez, A. K. Naik, F. Tobar et al. 2011. "The Human Genome: a Multifractal Analysis". *BioMed Central Genomics* 12(1):506.

Ouahabi, A. 2012. *Signal and Image Multiresolution Analysis*. John Wiley & Sons.

Pipiras, V., and M. S. Taqqu. 2017. *Long-Range Dependence and Self-Similarity*, Volume 45. Cambridge: Cambridge University Press.

Shekatkar, S. M., Y. Kotriwar, K. Harikrishnan, and G. Ambika. 2017. "Detecting Abnormality in Heart Dynamics from Multifractal Snalysis of ECG Signals". *Scientific Reports* 7(1):15127.

Shumway, R., and D. Stoffer. 2006. *Time Series Analysis and Its Applications*. 2nd ed. Switzerland: Springer Science and Business Media, LLC.

Thompson, J. R., and J. R. Wilson. 2016. "Multifractal Detrended Fluctuation Analysis: Practical Applications to Financial Time Series". *Mathematics and Computers in Simulation* 126:63–88.

Wendt, H. 2008. *Contributions of Wavelet Leaders and Bootstrap to Multifractal Analysis: Images, Estimation Performance, Dependence Structure and Vanishing Moments. Confidence Intervals and Hypothesis Tests.* Ph. D. thesis, Ecole normale supérieure de lyon-ENS LYON.

Wendt, H., P. Abry, and S. Jaffard. 2007. "Bootstrap for Empirical Multifractal Analysis". *Institute of Electrical and Electronics Engineers, Inc. Signal Processing Magazine* 24(4):38–48.

Wharton 2013. "Wharton Research Data Services: http://wrds-web.wharton.upenn.edu/wrds/".

AUTHOR BIOGRAPHIES

DAMON FREZZA is a Senior Operations Research Analyst at The MITRE Corporation. He holds a BS in Interdisciplinary Mathematics and Economics from Bucknell and an MS in Operations Research from Georgia Tech. His research interests include stochastic modeling and optimization. His email address is dfrezza@mitre.org.

JAMES R. THOMPSON is an Operations Research Principle at The MITRE Corporation. He holds a BS in Industrial and Systems Engineering from Georgia Tech and an MS and PhD in the same field from North Carolina State. His research interests include the implementation of stochastic modeling techniques and data analysis. His email address is j7sthompson@gmail.com.

DAVID SLATER is a Lead Operations Research Analyst at The MITRE Corporation. He holds a BS in Mathematics from Michigan Technological University and a MS and PhD in Applied Mathematics from Cornell University. His research interests include complex systems analysis, network science, and complexity science. His email address is dslater@mitre.org.

GARRY JACYNA is a MITRE Fellow and a recognized expert in signal processing, specializing in analytical-based performance analysis models for Navy and DHS programs, including novel applications of cutting-edge algorithms. Garry holds a B.S. in physics and an M.S. and Ph.D in Mathematics from from Rensselaer Polytechnic Institute. His email address is gjacyna@mitre.org.