Actuarial Applications of Multifractal Modeling Part II: Time Series Applications

Yakov Lantsman, Ph.D. and
John A. Major, ASA, MAAA

Actuarial Applications of Multifractal Modeling

Part II: Time Series Applications

by Yakov Lantsman, Ph.D. and John A. Major, ASA, MAAA

email:lantsman@netrisk.com, jmajor@guycarp.com

Abstract

Multifractals are mathematical generalizations of fractals, objects displaying "fractional dimension," "scale invariance," and "self-similarity." Many natural phenomena, including some of considerable interest to the casualty actuary (meteorological conditions, population distribution, financial time series), have been found to be well-represented by (random) multifractals. In part II of this paper, we show how to fit multifractal models in the context of one-dimensional time series. We also present original research on the multifractality of interest rate time series and the inadequacy of some state-of-the-art diffusion models in capturing that multifractality.

Introduction

In the accompanying part I paper, we introduced the ideas of fractal point sets and multifractal fields. We showed that those mathematical constructs are applicable to a wide range of natural phenomena, many of which are of considerable interest to the casualty actuary. We showed how to analyze sample data from multidimensional random fields, detect and measure multifractal behavior, fit a "universal" model, and use that model to simulate independent realizations from the underlying process. In particular, we discussed synthetic geocoding and the simulation of non-hurricane atmospheric perils.

The theory of self-similar random time series is more fully developed than the general multidimensional case. In this part II paper, we focus on time series analysis and financial applications. We present some additional theoretical machinery here and discuss applications to weather derivatives and financial modeling.

Time Series

Introduction to Multifractal Time Series Analysis; Structure Function

Financial and geophysical time series feature a large range of time scales and they are governed by strongly non-linear processes; this suggests the possible applicability of scaling (multifractal) models. We consider a random process $X(t)$ defined on the time segment $[0, T]$. The process $X(t)$ has variously represented exchange rates, interest rates, temperature and precipitation in our work.

As in the two-dimensional case, scale invariance is most readily tested by computing $P(k)$, the power spectrum of $X(t)$. In the case of a one-dimensional time series, standard techniques of spectral (Fourier) analysis are available in many off-the-shelf statistical and mathematical packages, including Microsoft EXCEL.

For a scaling process, one expects power law behavior:

$$P(k) \propto k^{-\beta} \quad (1)$$

over a large range of wave-numbers k (inverse of time). If $\beta < 1$, the process is stationary in the most accepted sense of the word [1], that is, $X(t)$ is statistically invariant by translation in t. If $1 < \beta < 3$, the process is non-stationary but has stationary increments and, in particular, the small-scale gradient (derivative or first difference) process will be stationary. Introducing the Hurst exponent H ($0 < H < 1$), a parameter describing the degree of stationarity of $X(t)$, we can express the exponent β as follows:

$$\beta = 2H + 1 \quad (2)$$

We can demonstrate a wide range of self-similar processes by changing the Hurst exponent: Brownian motion ($H = 0.5$, $\beta = 2$), an "anti-persistent" fractional Brownian motion ($0 < H < 0.5$, $1 < \beta < 2$), and "persistent" fractional Brownian motion ($0.5 < H < 1$, $2 < \beta < 3$). This is the class of additive models. The last has become popular for modeling financial time series.

Most of financial and geophysical time series demonstrate non-stationary behavior. This creates major complications if power spectrum analysis is the only available tool. It is well known [2] that knowledge of β alone is insufficient to distinguish radically different types of statistical behavior (the phenomena of "spectral ambiguity"). It is not so difficult to construct two processes with identical power spectra – one additive and sufficiently smooth, and the other one multiplicative with a high degree of intermittency. But such cases can be resolved with the help of multifractal analysis, which can be viewed as an extension in the time domain of scale-invariant spectral analysis.

An appealing statistical characteristic to use in exploring time series is the *structure function*. Structure function analysis of processes with stationary increments consists of studying the scaling behavior of non-overlapping fluctuations $\Delta X_\tau(t) = | X(t+\tau) - X(t) |$ for different time increments τ. One estimates the statistical moments of these fluctuations, which – assuming both scaling (1) and statistical translational invariance in time (i.e., the property of stationarity increments) – depend only on the time increment τ in a scaling way:

$$E(\Delta X_\tau(t)^q) \sim E(\Delta X_T{}^q)\left(\frac{\tau}{T}\right)^{\zeta(q)} \quad (3)$$

where $E(\Delta X_T{}^q)$ is a constant (T is the fixed largest time scale), $q > 0$ is the order of the moment, and $\zeta(q)$ is the scale invariant structure function. The expectation $E(\Delta X_\tau(t)^q)$ is assumed finite for q in an interval $[0, q_{max})$. The structure function $\zeta(q)$ is a focal concept in the one-dimensional theory of multifractals.

We examine some properties of $\zeta(q)$. By definition, we have $\zeta(0) = 0$. Davis A. et al. [1] show that $\zeta(q)$ will be concave: $d^2\zeta(q)/dq^2 < 0$. This is sufficient to define a "hierarchy of exponents" using $\zeta(q)$:

$$H(q) = \frac{\zeta(q)}{q} \tag{4}$$

It can also be shown that $H(q)$ is a non-increasing function. The second moment is linked to the exponent β as follows:

$$\beta = 1 + \zeta(2) = 2H(2) + 1 \tag{5}$$

Obtaining $\zeta(q)$ or, equivalently, $H(q)$ is the goal of structure function analysis. A process with a constant $H(q)$ function could be classified as "monofractal" or "monoaffine"; in the case of decreasing $H(q)$, multifractal or "multiaffine."

Additive processes can be shown to have linear $\zeta(q)$ or constant $H(q)$. For Brownian motion we have:

$$\zeta_{BM}(q) = \frac{q}{2} \tag{6}$$

For fractional Brownian motion (the fractional integration of order h of a Gaussian noise):

$$\zeta_{FrBM}(q) = q(h - \tfrac{1}{2}) \tag{7}$$

Note that Brownian motion corresponds to $h = 1$ (an ordinary integral of Gaussian white noise, which gives $H = ½$ in Fourier space).

In the case of the more exotic "Lévy flight" (additive processes with Lévy noise) the behavior of $\zeta(q)$ is still linear. In this case, there is a Lévy index α ($0 \le \alpha \le 2$), which characterizes the divergence of the moments of the Lévy noise. In general $\zeta(q)$ diverges for $q > \alpha$, but for finite samples we obtain the following $\zeta(q)$ function for a Lévy flight of index α:

$$\zeta_{LM}(q) = \frac{q}{\alpha} \tag{8}$$

for $q < \alpha$, and $\zeta(q) = 1$ for $q \ge \alpha$.

We thus see that observing non-linearity of an empirical $\zeta(q)$ function is a solid argument against the validity of an additive model. Below, we will show strong signs of curvature in the behavior of some empirical $\zeta(q)$ functions for financial time series.

The generic multifractal processes (non-linear, non-additive) could be modeled by multiplicative cascades. The central part of a multiplicative cascades is the generator (MCG, discussed in the part I paper) which should be represented by some infinitely divisible probability distribution. Using "canonical representation" (the Lévy-Khinchine representation) for infinitely divisible random variables, and arguments similar to those for the $K(q)$ function for the general D-dimensional case, we obtain the following "universal form" for the structure function of a non-stationary process:

$$\zeta(q) = qH - \frac{C1}{\alpha - 1}(q^{\alpha} - q) \tag{9}$$

where $H = \zeta(1)$ the same as (2), $C1$ is a parameter with the same role as in equation (31) of part I, and α is the Lévy index.

Analogous considerations could guide us to modify part I's equation (31) to express the $K(q)$ function for a non-conservative field:

$$K(q) + qH = \begin{cases} \dfrac{C1}{\alpha - 1}(q^{\alpha} - q) & \alpha \neq 1 \\ C1 q \log(q) & \alpha = 1 \end{cases} \tag{10}$$

where the H parameter is the degree of non-stationarity of the process. In other words, first bring the field to a state of stationarity (by fractional differentiation, i.e., power-law filtering in Fourier space or a small-scale gradient transformation) to eliminate the linear part qH, and then proceed with the analysis as for conservative fields.

To summarize, the basic steps are:

1. Examine the data for evidence of intermittency and self-similarity; this could be accomplished by studying the power spectrum.
2. Establish the status of multifractality (or monofractality) and qualitatively characterize the system under investigation; for this, we use the structure function.
3. Fit model parameters to the universal form of $\zeta(q)$.
4. Simulate, using multiplicative cascade techniques based on the universal form of the generator.
5. Apply, including, possibly, drawing inferences about the underlying process.

A Growing Crisis in Financial Time Series Modeling?

There is a growing awareness among researchers that the existing "classical" models cannot accommodate some essential properties underlying financial phenomena. The accumulation of a tremendous amount of highly reliable data from the financial markets around the world reveals distinctive characteristics of financial time series that had previously been overlooked because of lack of data. Some of the most important features are:

- scaling or self-similarity (at different time scales);
- long-term memory or persistency;
- volatility clustering;
- hyperbolic or "Paretian" tails.

To compensate for the consequences of these characteristics, the number of parameters in the "classical" models has been increasing over time. If this continues unchecked, it could make models unstable and decrease their predictive power.[1]

We distinguish two major classes of models in use by practitioners today: continuous time stochastic diffusion models ("diffusion models") and discrete time series models ("discrete models").

Diffusion models build on the well-understood theory of Brownian motion. The development of stochastic calculus (particular Itô integrals) and the theory of martingales created the essential mathematical apparatus for equilibrium theory. The assumption of arbitrage free pricing (rule of one price) has a very elegant mathematical interpretation as a change of stochastic measure and the transformation to a risk-neutral stochastic process.

Application of diffusion models is a crucial element in the valuation of a wide variety of financial instruments (derivatives, swaps, structured products, etc). Researchers have, however, long recognized major discrepancies between models based on Brownian motion and actual financial data, including long-term memory, volatility clustering and fat tails. To resolve these problems some extensions of diffusion models were offered. Often, this means introducing more stochastic factors, creating so-called multi-factor models.

Modern discrete models extend classical auto-regressive (AR) moving average (MA) models with recent advances in the parameterization of time-conditional density functions. These include ARCH, GARCH, PGARCH, etc. Discrete models have been partially successful in compensating for lack of long-term memory, volatility clustering and fat tails, but at the cost of an increasing number of parameters and structural equations. Using appropriate diagnostic techniques one can demonstrate that the statistical properties of discrete models (viz., self-similarity of moments, long-term memory, etc.) are essentially the same as for Brownian motion.

There is a third class of models, in little use by practitioners, but familiar to academics. This group constitutes the so-called additive models, including fractional Brownian motion, Lévy flight and truncated Lévy flight models. These models can replicate *mono*-fractal structure of underlying processes – their corresponding structure functions $\zeta(q)$ (7), (8) are linear – but they cannot produce *multi*fractal (nonlinear) behavior.

Case Study: Foreign Exchange

Here, we present an example of the application of multifractal analysis to exchange rate modeling, substantially following Schmidt, F. [3]. Figure 1 represents the exchange rate time series (US$/GDM spot rate 1975 - 1990 weekly observations) and Figure 2 the corresponding logarithmic changes in exchange rate.

Figure 3 represents a power spectrum analysis (in log-log space) of the FX time series. Visual inspection, and the close fit of the regression line, supports the hypothesis of scale-invariant behavior. The power spectrum obeys a power law (Equation 1). The slope

[1] A similar "Ptolemaic crisis" afflicted meteorological precipitation modeling in the 1980s. See, e.g., the Water Resources Research special issue on Mesoscale Precipitation Fields, August 1985.

of the straight line is the parameter β; here equal to 1.592. This value suggests the underlying process may be non-stationary but with stationary increments.

An important application of multifractal analysis is to characterize *all* order moments for the validation of a scaling model. The appropriate tool to do this for the particular case of a time series is structure function analysis.

To apply the structure function method, we rewrite the equation (3) in logarithmic form:

$$\log[E(\Delta X_\tau(t)^q)] = \log[E(\Delta X_T{}^q)] + \zeta(q)\{\log(\tau) - \log(T)\} \tag{11}$$

The expectation $E(\Delta X_\tau(t)^q)$ is estimated by the so-called *partition function*

$$E(\Delta X_\tau(t)^q) \cong \frac{1}{N}\sum_\tau |\Delta X_\tau(t)|^q \tag{12}$$

(see Fisher, A. et al. [4]). We then plot $\log[E(\Delta X_\tau(t)^q)]$ against $\log(\tau)$ for various values of q and various values of τ. Linearity of these plots for given values of q indicates self-similarity. Linearity could be checked by visual inspection or by some more sophisticated techniques (e.g., significance test for higher-order regression terms). The slope of the line, estimated by least squares regression, gives an estimate of the scaling function $\zeta(q)$ for that particular q.

The structure function, mapping q to its slope, is depicted in Figure 4. Here, we also draw an envelope of two straight lines corresponding to Brownian motion (slope 0.5) and fractional Brownian motion (slope 0.6), respectively. The non-linear shape of the empirical curve is the signature of multifractality.

Having established the existence of multifractality in the data, we can move to the next step – fitting parameters. In the case of one dimensional (time series) field, we use equation (9) to find universal parameters. For FX data, the universal parameters are: H = 0.532, α = 1.985, $C1$ = 0.035.

Case Study: Interest Rates

In this section, we present an original analysis of US interest rates. We use weekly observations of 3-month Treasury Bill Yield Rates (1/5/1962 - 3/31/1995). Figure 5 represents the interest rate time series and Figure 6 the corresponding logarithmic changes in interest rate from one period to the next.

We start with the power spectrum in Figure 7. Visual inspection and regression confirm the hypothesis of scaling behavior with corresponding β_{eir} = 1.893. The value of β_{eir} indicates that this interest rate series *might* be modeled by a non-stationary process with stationary increments.

Figure 8 represents the $\zeta(q)$ curve for interest rates with the same Brownian motion and fractional Brownian motion lines that we used for the FX analysis overlaid on the graph for reference. Again, the signature of multifractality is clearly present in the data. We obtain the following universal values: H = 0.612, α = 1.492, $C1$ = 0.095. These values could be used to simulate interest rates by applying a multiplicative cascade technique.

Andersen-Lund is Not Multifractal

We present an original analysis of the three-factor Anderson-Lund model of interest rates and show that even this model, with its highly complex structural equations and difficult fitting techniques, cannot replicate key features of empirical interest rate data.

The general form of the diffusion model Vetzal, K. [5] is

$$dX = \mu(X,t)\cdot dt + \sigma(X,t)\cdot dW \tag{13}$$

where X is the (possibly vector) random variable evolving over time, μ is a (possibly vector) function describing the instantaneous rate of change of X at a point in time, and σ is a (possibly matrix) function describing the instantaneous impact of changes in the (possibly vector) Gaussian random walk W, i.e. dW is a (are independent) Gaussian random variable(s). In the multidimensional case, the dimension of X does not necessarily equal the dimension of W. For interest rate models, one of the elements of X will represent the short rate of interest.

The primary purpose of these models has been to develop arbitrage-free prices of illiquid bonds, interest rate derivatives, etc., so the primary consideration has been fidelity in reproducing available market data, in particular, yield curves. However, obtaining realistic depictions of the objective behavior of the short rate (the historical evolution of the short rate over time) has been an important secondary consideration. It is this tension between the desire for analytical tractability on the one hand and realism on the other that has driven the development of ever more sophisticated models.

The simplest such model is Merton [6]:

$$dr_t = \theta\cdot dt + \sigma\cdot dW \tag{14}$$

where r_t is the interest rate at time t, θ is the average growth rate of the process, and σ is a volatility scale parameter.

Perhaps the most sophisticated of the analytically tractable models is the Cox-Ingersoll-Ross (CIR) model [7]:

$$dr_t = \kappa\cdot(\theta - r_t)dt + \sigma\cdot r_t^{1/2}\cdot dW \tag{15}$$

where κ is the mean reversion constant and θ is the global mean of the process. CIR adds realism to the Merton model by introducing mean reversion and volatility that is functionally dependent on the level of the rate.

Visual analysis of the interest rate time series graphs (as well as statistical diagnostics) reveals several distinctive features to US interest rates that cannot be accommodated by the CIR model.

1. Local trends in interest rate movements, indicating a changing mean to which the process reverts.
2. Heteroscedasticity that is not simply a function of the level of the rates.
3. Volatility clustering.

To address these limitations of CIR and previous models, Andersen and Lund [8] introduced the following (analytically intractable) three-factor model:

$$dr_t = \kappa_1 \cdot (\mu_t - r_t)dt + \sigma_t \cdot r_t^{\gamma} \cdot dW_{1,t} \tag{16}$$

$$d\log\sigma_t^2 = \kappa_2 \cdot (\alpha - \log\sigma_t^2)dt + \xi_1 \cdot dW_{2,t} \tag{17}$$

$$d\mu_t = \kappa_3 \cdot (\theta - \mu_t)dt + \xi_2 \cdot \mu_t^{1/2} \cdot dW_{3,t} \tag{18}$$

where r_t is the interest rate at time t, σ_t is the (unobserved) volatility, μ_t is the (unobserved) local mean of the process, k_1, k_2, k_3 are mean reversion constants, θ is the global mean of the process, α is the global mean of the log-volatility process, γ, ξ_1, ξ_2 are parameters, and $W_{1,t}$, $W_{2,t}$, $W_{3,t}$ are independent Gaussian random variables. Equation (16) can be seen as a generalization of the CIR model with unknown parameter γ instead of ½, and the fixed mean and volatility terms being replaced by endogenous variables evolving through their own diffusions (Equations 17 and 18).

The Andersen-Lund model represents the most realistic diffusion model we are currently aware of. The price of its realism is the need for substantial computing power. To fit the parameters, Andersen and Lund use the so-called Efficient Method of Moments procedure Gallant and Tauchen [9], which is an iterative method involving many simulations of the diffusion process. Calculation of yield curves similarly requires many simulation cycles, as there is no (known) closed-form solution.

Just how realistic is it? Figure 9 shows a 5,000-quarter simulation of interest rates using the A-L model with their recommended parameters. Figure 10 is the corresponding logarithmic changes in interest rate. A visual comparison of these graphs with the corresponding empirical interest rate graphs, and a cursory statistical examination of same, seems to validate the A-L approach to modeling interest rate time series. We will demonstrate that a deeper analysis of the scaling properties of *all* moments (not just the first and second) reveals fundamental differences between the A-L simulation and the empirical data. The simulated A-L data does not exhibit the multifractality that real interest rates possess.

Figure 11 shows the power spectrum function of the simulated time series out of the A-L model. Here, the parameter $\beta_{sir} = 1.772$, which is fairly close to the value obtained for the empirical data. The power spectrum, however, represents second moment statistics only. Its slope is not sufficient to validate a particular scaling model: it gives only partial information about the statistics of the process. One would need full knowledge of the probability distribution of the process or, equivalently, all of its statistical moments (not just second order) for a full validation.

Figure 12 represents the $\zeta(q)$ curve for the A-L simulated interest rate series, with the usual Brownian motion and fractional Brownian motion lines drawn for reference. Visual inspection and statistical testing indicate that the structure function of the data simulated by the A-L model and that of Brownian motion are nearly identical; the stochastic process underlying the A-L model appears to be *mono*fractal.[2]

The fundamental difference in scaling behavior revealed by the structure function comparison could lead to qualitatively different time series behavior. The universal parameters fit to the empirical process in the previous section indicate that the underlying mechanism should have a multiplicative cascade structure with (approximate) Lévy generator, rather than an additive process of information accumulation (Brownian motion type). Paraphrasing Müller et al. [10], the large scale volatility predicts small scale volatility much better than the other way around. This behavior can be compared to the energy flux in hydrodynamic turbulence, which cascades from large scales to smaller ones, not vice-versa.

Conclusions

In the companion part I paper, we introduced the ideas of fractal point sets and multifractal fields. We showed that while those mathematical constructs are rather bizarre from a traditional point of view (e.g., theory of smooth, differentiable functions), they nonetheless have applicability to a wide range of natural phenomena, many of which are of considerable interest to the casualty actuary. We showed how to analyze sample data from multidimensional random fields, detect scaling through the use of the power spectrum, detect and measure multifractal behavior by the trace moments and double trace moments techniques, fit a "universal" model to the trace moments function K(q), and use that model to simulate independent realizations from the underlying process by a multiplicative cascade. In particular, we discussed synthetic geocoding and the simulation of hail and tornadoes.

In this part II paper, we showed how to analyze time series through the structure function, and showed particular examples of foreign exchange and interest rate time series. We discussed the variety of time series models in use by practitioners and theoreticians and showed how even state-of-the-art diffusion models are not able to adequately reflect the multifractal behavior of real financial time series.

The field of stochastic modeling is constantly growing and evolving, so the term "Copernican revolution" might be too strong to describe the advent of multiplicative cascade modeling. Nonetheless, multifractals have clearly taken hold in the realm of geophysical and meteorological modeling, and it seems clear that they will eventually find their place in the world of financial models, as well. However, there are still numerous open questions, such as how to implement arbitrage-free pricing, that need to be answered before multifractal models can replace diffusion models as explanations of market pricing mechanisms.

[2] Theoretical arguments suggest monofractality for any additive models, Schmidt [3].

References

1. A. Davis, A. Marshak, W. Wiscombe, and R. Cahalan, "Multifractal characterizations of nonstationarity and intermittency in geophysical fields: Observed, retrieved, or simulated," *Journal of geophysical research*, Vol. 99, N. D4, pp.8055-8072, April 20, 1994.
2. D. Schertzer, and S. Lovejoy, "Physical modeling and analysis of rain clouds by anisotropic scaling multiplicative processes," *Journal of geophysical research*, Vol. 92, pp. 9693-9714, 1987.
3. F. Schmitt, D. Schertzer, and S. Lovejoy, "Multifractal analysis of foreign exchange data," submitted to *Applied Stochastic Models and Data Analysis (ASMDA)*.
4. A. Fisher, L. Calvet, and B. Mandelbrot, "Multifractality of deutschmark / US dollar exchange rates," *Cowles Foundation Discussion Paper* # 1165, 1997.
5. K. Vetzal, "A survey of stochastic continuous time models of the term structure of interest rates," *Insurance: Mathematics and Economics* #14, pp. 139-161, 1994.
6. R. C. Merton, "Theory of rational option pricing," *Bell Journal of Economics and Management Science*, Vol. 4, pp. 141-183, 1973.
7. J. Cox, J. Ingersoll, and S. Ross, "A theory of the term structure of interest rates," *Econometrica* # 53, pp. 385-407, 1985.
8. T. Andersen, and J. Lund, Stochastic "Volatility and mean drift in the short rate diffusion: sources of steepness, level and curvature in the yield curve," *Working Paper* #214, 1996.
9. A. Gallant, and G. Tauchen, "Estimation of continuous time models for stock returns and interest rates," *Manuscript*, Duke University, 1995.
10. U. Müller, M. Dacorogna, R. Dave, R. Olsen, O. Pictet, and J. von Weizsacker, "Volatilities of different time resolutions – Analyzing the dynamics of market components," *J. of Empirical Finance* # 4, pp. 213-239, 1997.

Figures for Part II

Yakov Lantsman

John A. Major

Figure 1: US$/GDM Exchange rate time series

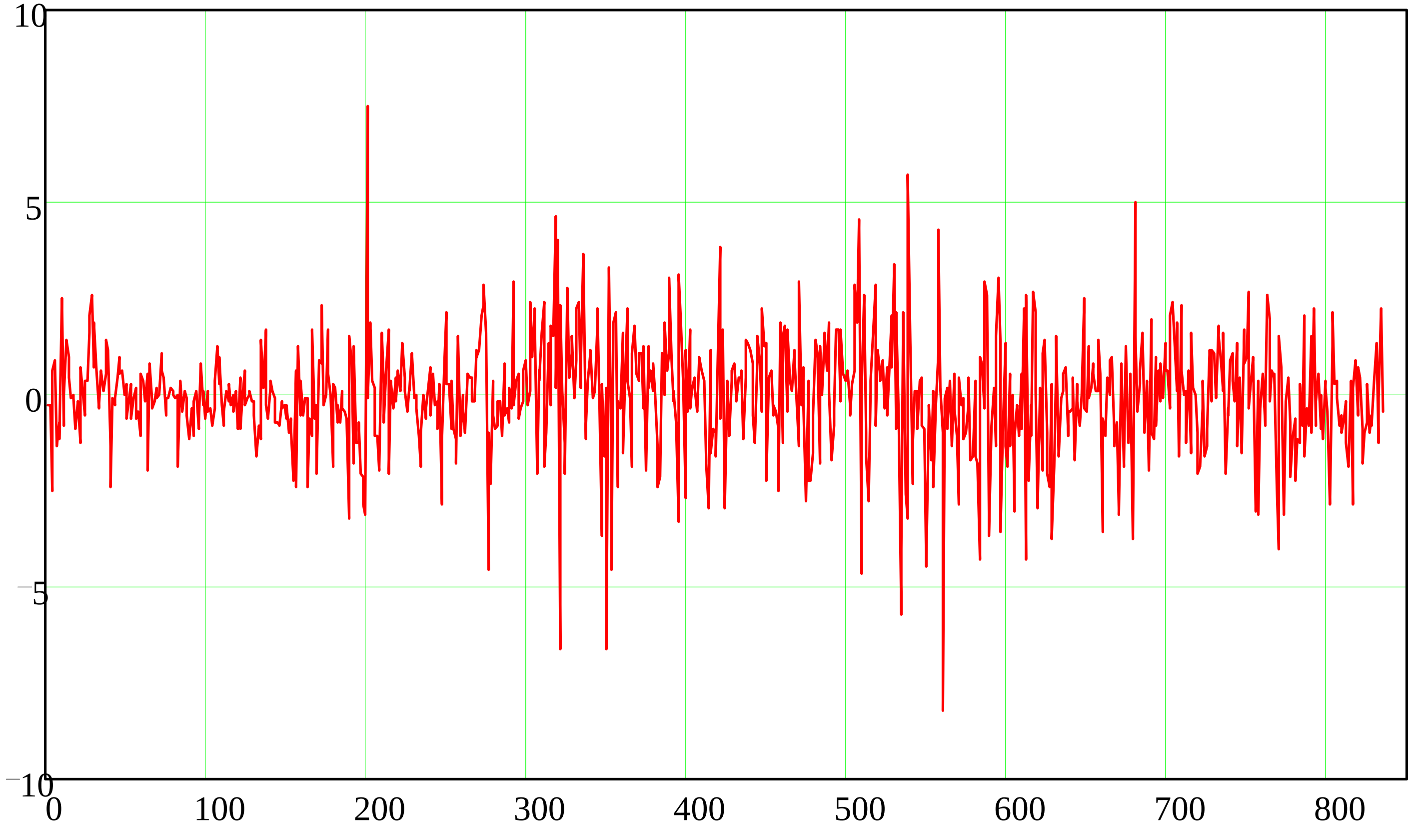

Figure 2: Logarithmic changes of Figure 1

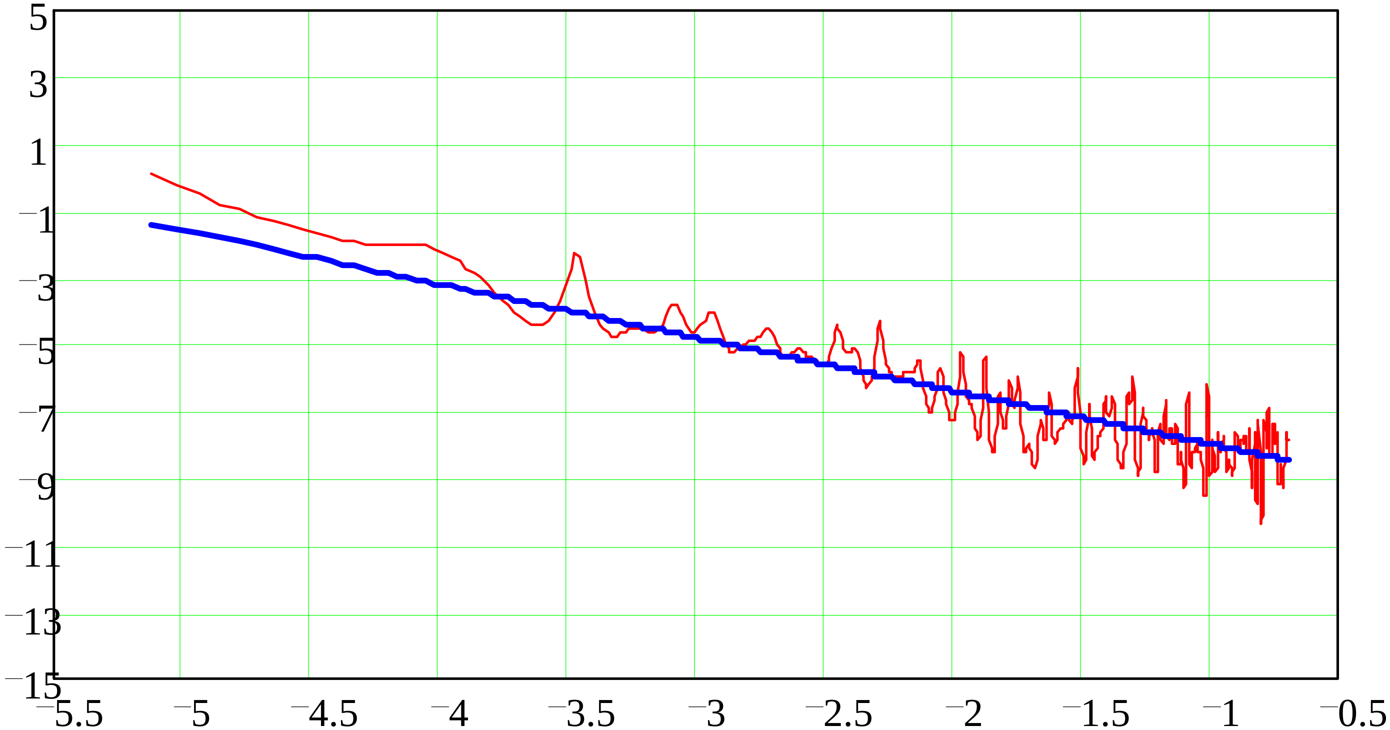

Figure 3: Power spectrum of FX data (log-log)

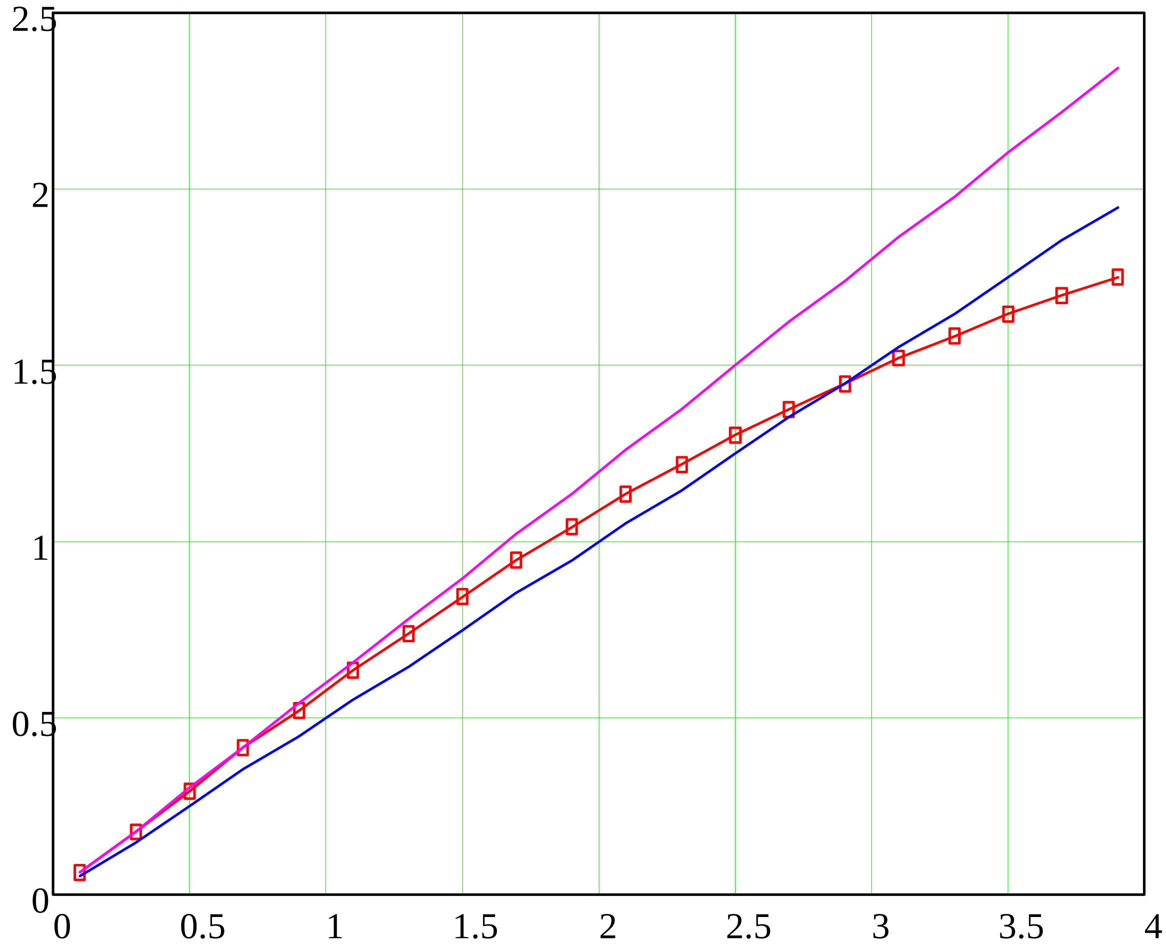

Figure 4: Structure Function Curve for FX

Figure 5: 3-mo T-Bill rates

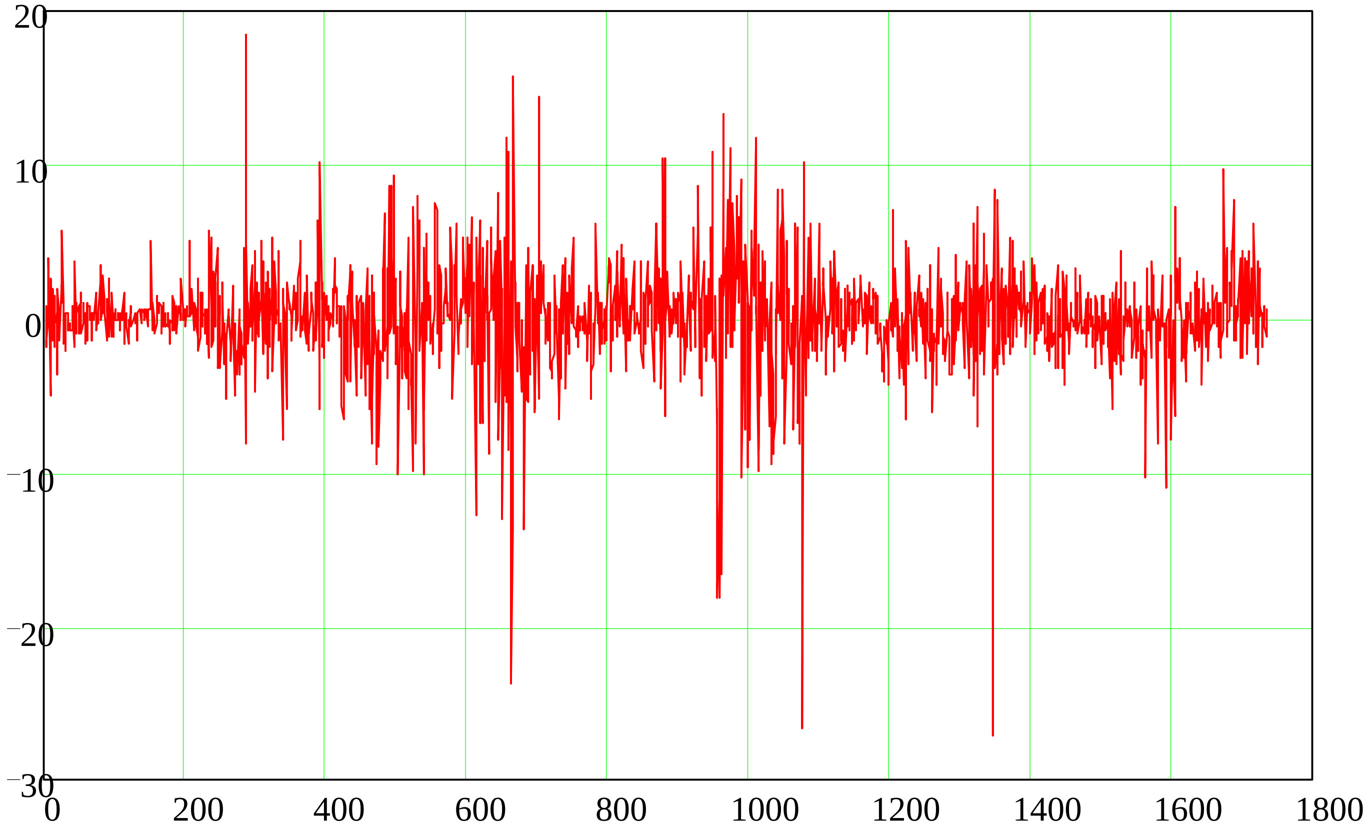

Figure 6: Logarithmic changes of Figure 5

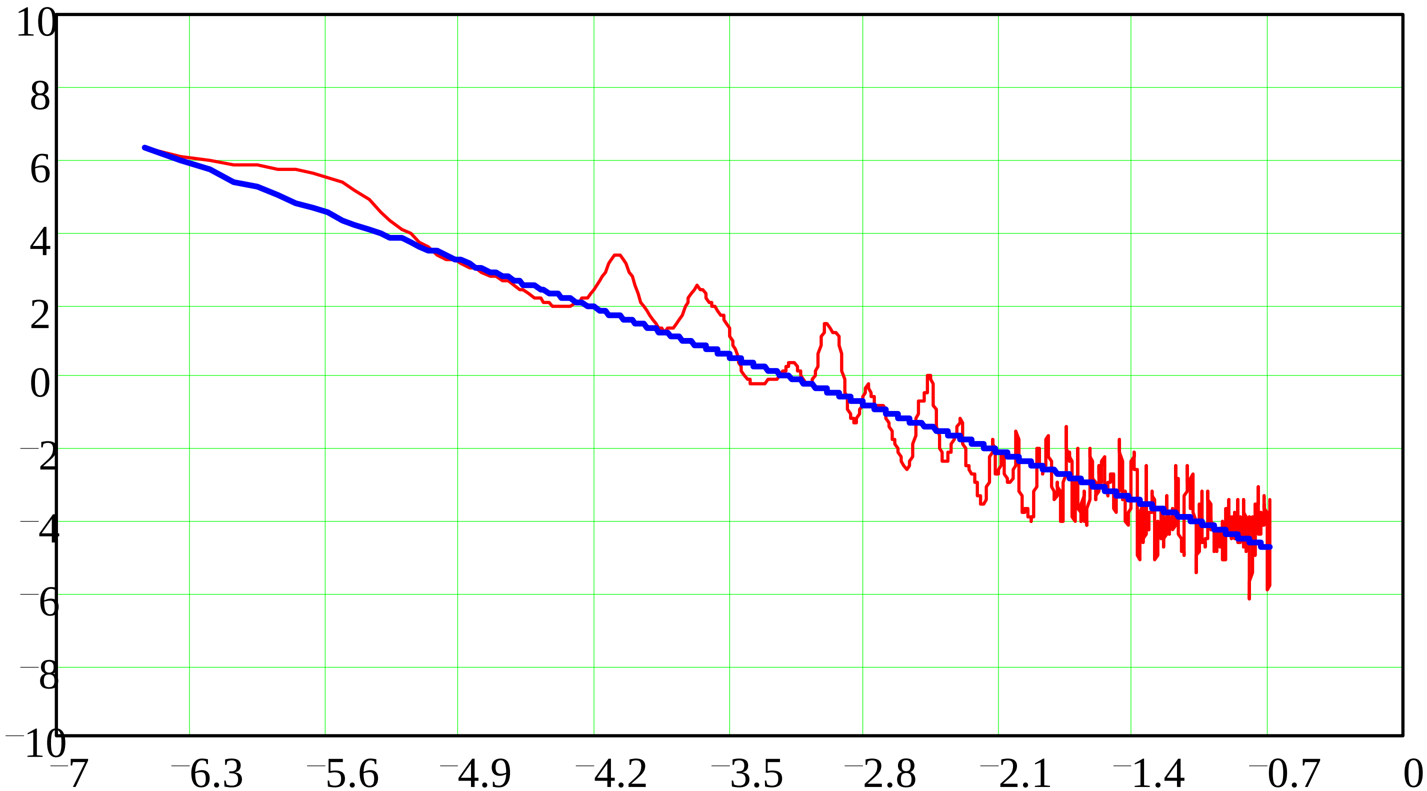

Figure 7: Power spectrum of interest rate

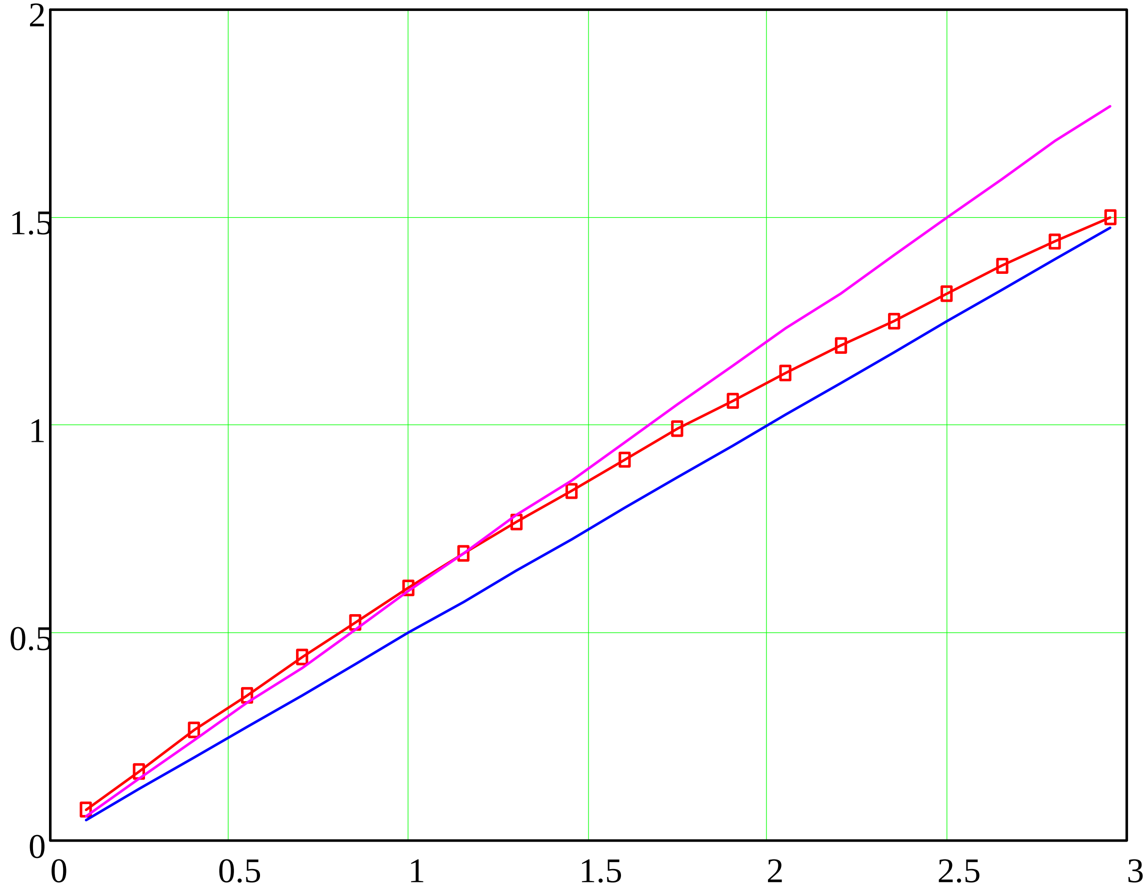

Figure 8: Structure Function for Interest Rate

Figure 9: 5,000 simulated interest rates

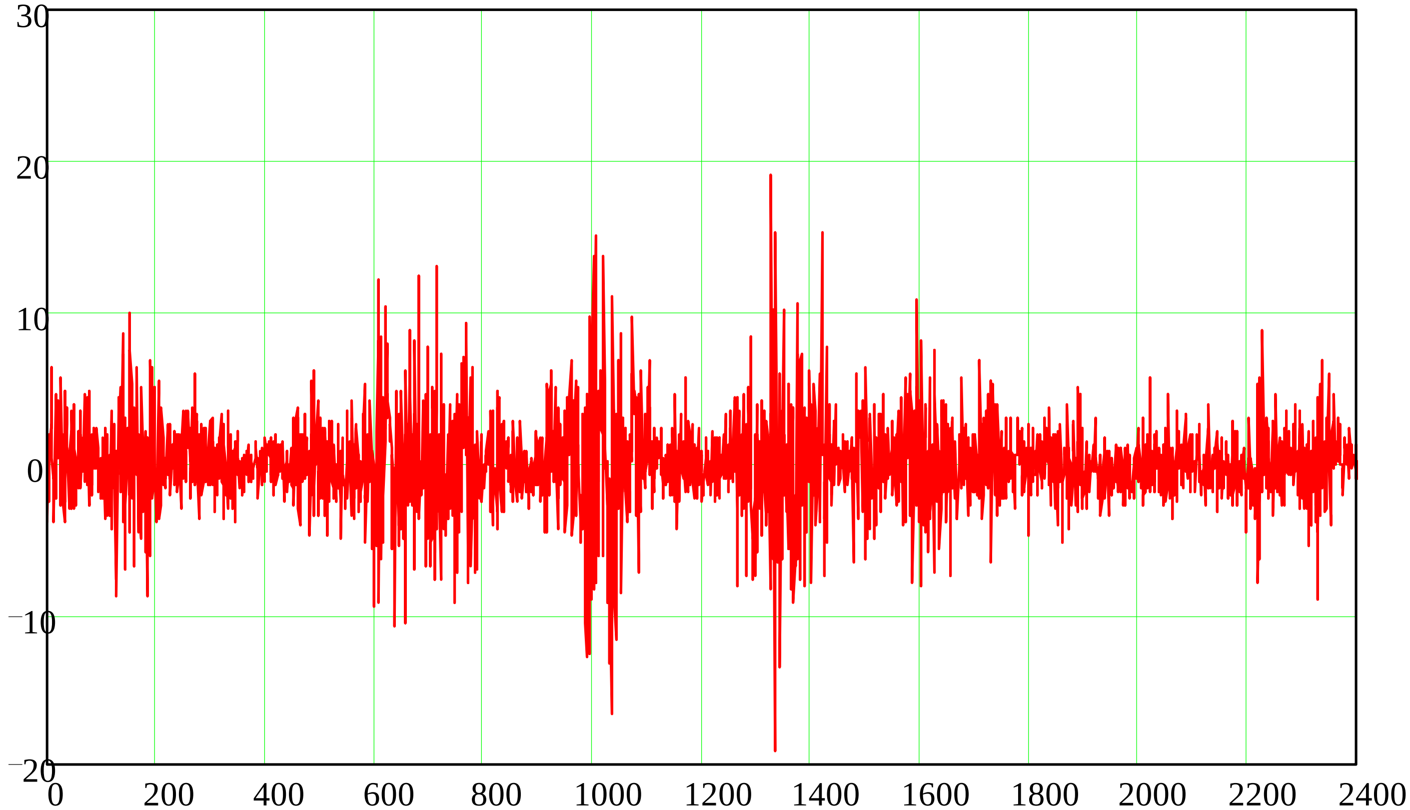

Figure 10: Logarithmic changes of Figure 9

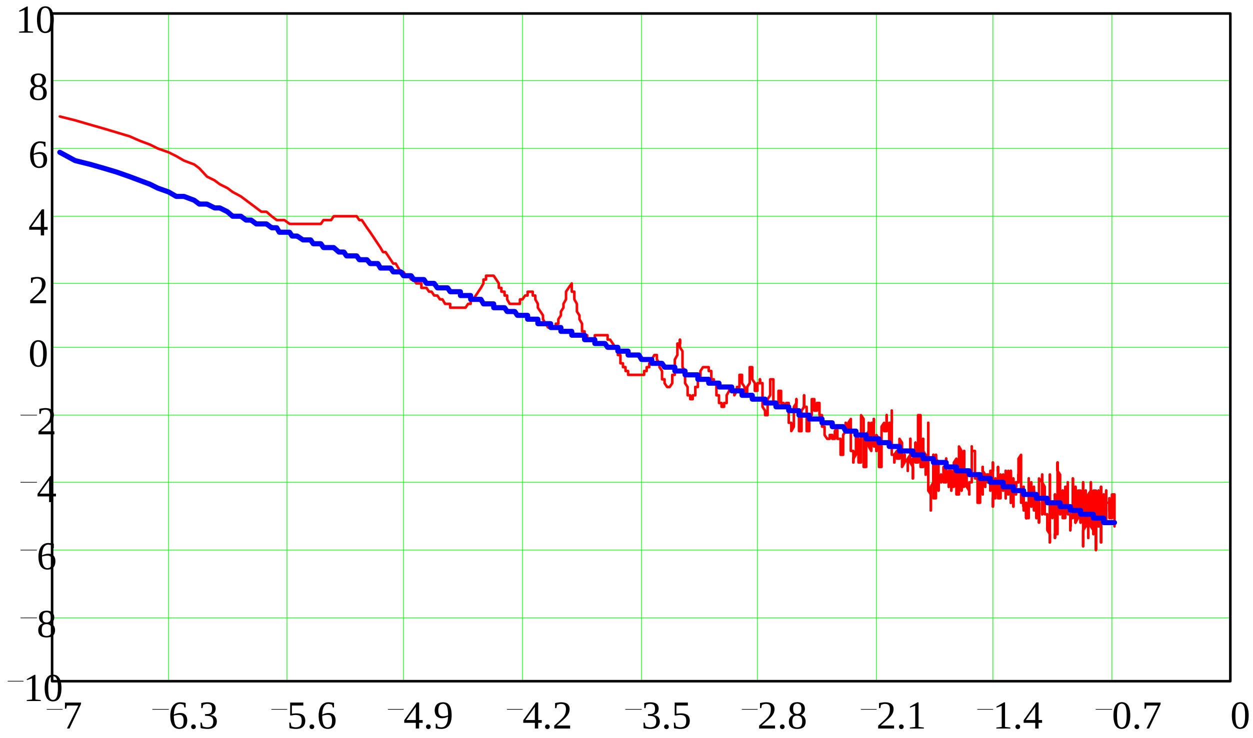

Figure 11: Power spectrum of simulated interest rate

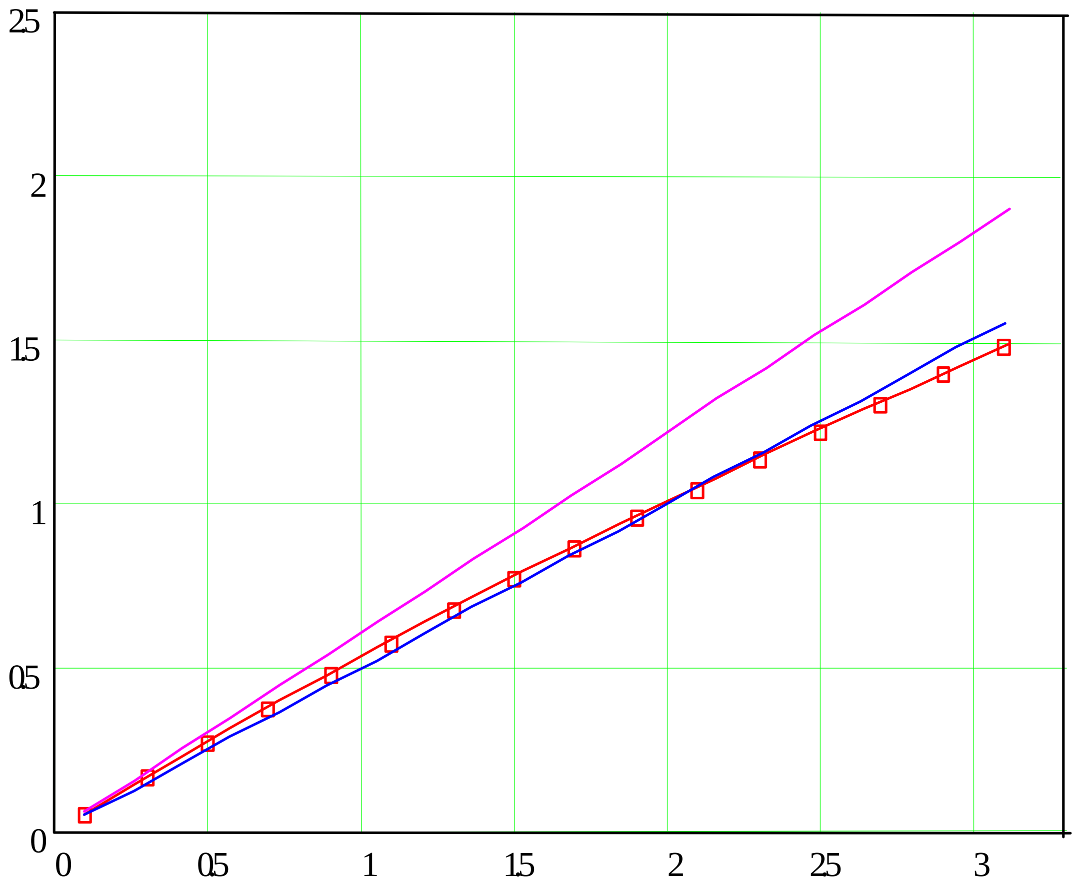

Figure 12: Structure Function for simulated interest rate