

TESTING FOR NON-LINEAR STRUCTURES IN ARTIFICIAL AND REAL-WORLD FINANCIAL DATA WITH RECURRENCE QUANTIFICATION ANALYSIS

Ricardo Giglio, Eduardo Ferreira da Silva 

Department of Systems and Production Engineering, Federal University of Santa Catarina.
Campus Reitor João David Ferreira Lima, s/n - Trindade, Florianópolis - SC, 88040-900, Brazil
{ricardo.giglio, e.f.silva}@ufsc.br

Abstract – Early applications of empirical methods from chaos theory suggested the existence of low dimensional chaotic motion in empirical financial data. However, such results were questioned, and it is then believed that the search for low dimensional chaos in financial data was not successful. On the other hand, at the same time that the hypotheses that raw returns are independent and identically distributed (*IID*) is often rejected, they indeed present a quite small degree of autocorrelation. These facts suggest that prices in financial markets do not behave completely at random, although their hidden structures seem more complex than those observed in low dimensional chaotic systems. Previous work tested for non-linearity or the presence of low dimensional chaos in artificial financial data generated from the Lux-Marchesi model by means of the BDS and Kaplan tests. Addressing the same model, researchers extended those results by applying Hinich's bi-spectral and White's tests and introducing the application of Recurrence Quantification Analysis (*RQA*) on artificial financial data based on Recurrence Rate, Determinism, Entropy, and Maximal Diagonal Length. Contributing to this research, the present paper has two main goals: (i) to contrast previous findings with an *RQA* application on data generated by a more evolved of microscopic model of financial markets - the Structural Stochastic Volatility (*SSV*) model; and (ii) to extend the *RQA* investigation above with additional recurrence measures (namely, Divergence, Laminarity, and Maximal Vertical Length) being applied to distinct real-world financial data. The objective is to assess if *RQA* results could help to distinguish between artificial and real-world data, even if linearity is rejected in both cases. It is shown evidence, in agreement previous findings, to support the rejection of linearity or low dimensional chaotic motion in an artificial financial time series generated from the *SSV* microscopic model. In addition, it is also shown that that *RQA* measures can help to discriminate artificial from real-world financial data, at least when specific *RQA* measures are considered.

Keywords – Recurrence Quantification Analysis, Nonlinearity test, Artificial Financial Data

1 Introduction

According to Chen et al. [2], early applications of empirical methods from chaos theory suggested the existence of low dimensional chaotic motion in empirical financial data. However, such results were questioned, and it is then believed that the search for low dimensional chaos in financial data was not successful.

On the other hand, the independent and identically distributed (*IID*) hypotheses is often rejected, at the same time that raw returns present a quite small degree of autocorrelation. These facts suggest that prices in financial markets do not behave completely at random, although their hidden structures seem more complex than those observed in low dimensional chaotic systems.

An artificial financial data generator should be able to reproduce several of the stylized facts often observed in empirical financial data. For instance, the presence of a unit root in the price dynamics, the existence of heteroscedasticity (the so-called volatility clustering) and the heavy tails in the distributions of returns, and the long-term dependency of absolute returns. For an extensive list see [3]. Along with all these structures usually found in empirical data is the very low autocorrelation of returns which makes the financial markets appear efficient, at least by considering the weak form of information efficiency from the Efficient Market Hypothesis [10].

By considering artificial data generated from the Lux-Marchesi microscopic model [9], Chen et al. [2] and Belaire-Franch [1] state that the resulting price dynamics appear less random than the pseudo-random numbers used as shocks to the fundamental price of the model. Such results are supported by the empirical chaos literature, where randomly reshuffled series usually lead to higher estimates of the so-called correlation dimension.

Chen et al. [2] tested for non-linearity or the presence of low dimensional chaos in artificial financial data generated from the Lux-Marchesi model by means of the BDS and Kaplan tests. Belaire-Franch [1], also addressing the Lux-Marchesi model, extended those results by applying Hinich's bi-spectral and White's tests, and introducing the application of Recurrence Quantification Analysis (*RQA*) [22] on artificial financial data. Hence, the key results have shown that both simulated and empirical time series show traces of hidden structure, but apparently a more complicated one than that generated by low-dimensional deterministic dynamics.

Contributing to this research, the present paper has two main goals. The first is to compare such previous results with an application of the (*RQA*) methodology using data generated by a more evolved microscopic model, the Structural Stochastic

Volatility (*SSV*) model [6]. It is shown evidence in agreement with Chen et al. [2] and Belaire-Franch [1] to support the rejection of linearity or low dimensional chaotic motion in an artificial financial time series generated from the *SSV* microscopic model.

The second goal concerns the application of the same methodology in two contrasting real-world scenarios, in order to assess if *RQA* results could help to distinguish between artificial and actual financial data, even if linearity is rejected in both cases. In this case, results suggest that the *RQA* methodology can help to tell artificial and real-world financial data apart, at least when specific *RQA* measures are considered. This paper is based on the research conducted by one of the authors during his PhD [7].

In the next section, a simple chaotic system (the logistic map) is used as an illustration to highlight the properties of the Recurrence Quantification Analysis (*RQA*) based on the Recurrence Plots (*RP*) [5] for detecting linearity or chaos in time series. Then, surrogate non-linearity tests based on the *RQA* measures are briefly discussed.

The third section introduces three time series used in this paper. First, the microscopic generating mechanism of the Structural Stochastic Volatility (*SSV*) model [6] is briefly described. Afterward, two real-world financial time series with clearly distinct behaviors are presented, namely, daily returns from Bitcoin and from the S&P500 index.

Section four shows the results of two analysis performed using surrogate non-linearity tests based on *RQA*: one following the exact approach as in Belaire-Franch [1] using only *SSV* artificial data, and a second making the methodology more comparable among different time series using also real-world data. Finally, section five presents conclusions.

2 Recurrence Quantification Analysis

Recurrence Quantification Analysis (*RQA*) is a nonlinear method of analyzing dynamical systems originally proposed by Zbilut and Webber [22]. It is carried out by calculating some measures on the so-called Recurrence Plot (*RP*), which is a graphical representation of how often in time a trajectory visits the neighbor regions of its phase space.

This section is divided as follows: first, the Logistic Map [13] is introduced and used as an example for outlining the required procedures in the creation of an *RP*. Then, selected *RQA* complexity measures are calculated for different parameters of the logistic map in order to illustrate their properties. The last part concludes by presenting a surrogate non-linearity test based on *RQA* measures, which are used in next section to test artificial financial data generated by the *SSV* model.

2.1 Logistic Map Example

The logistic map is an early example of an apparently simple mathematical model presenting very complicated dynamics [13]. It consists of a polynomial mapping of degree 2 in the following form:

$$x_{n+1} = rx_n(1 - x_n) \quad (1)$$

where x_n is defined in $[0, 1]$ and can be interpreted as the ratio between some current population at time n and the environment total capacity, which is assumed to be fixed. The initial population ratio x_0 has also to be assumed fixed. The growth parameter r can be any positive number, although the population ratio interpretation only holds for $r \leq 4$. For higher values of r , x does not respect the $[0, 1]$ range. Values of r below 3 always present a fixed point as long term value, while for values higher than 3 periodic orbits of increasing order start to appear. Finally, for values higher than about 3.57 the logistic map exhibits chaotic behavior for most values beyond this threshold (except for some regions of periodicity such as $r = 3.83$).

2.2 Recurrence Plots

Recurrence plots (*RP*) were proposed by Eckmann et al. [4] as a tool for visualizing recurrences of dynamical systems. In order to build an *RP* for a given time series, one has to first reconstruct its phase-space trajectory. As stated by Takens' theorem [18], a topologically equivalent representation of the original high-dimension system can be reconstructed from a time sequence of observations of its states by means of the method of time delay. This procedure consists of generating, from the original time series X_t , the set of embedded vectors $X_t^m = X_t, X_{t+\tau}, X_{t+2\tau}, \dots, X_{t+(m-1)\tau}$, where m is the underlying embedding dimension of the system, and τ is the time delay used in the reconstruction of the phase space from a time series.

Among several methods for choosing an appropriate value for the embedding dimension m , Zbilut et al. [24] point out that a heuristic approach is well suited when the final objective is to generate an *RP*. This approach consists of choosing a very high value for m ($m > 20$), and of sequentially decreasing its value until significant differences in the *RP* appear. Let m^* denote the value of m at this point. They state that this difference is due to the existence of false nearest neighbors, and, thus, a value of m a few dimensions higher than m^* should be enough for embedding.

Having reconstructed the phase space, the objective of an *RP* is to check for recurrence patterns, that is, to verify whether the system roughly returns to neighborhoods already visited in the past. To do so, how close two trajectories should be in order to be considered a recurrence has also to be defined. This is done by setting a critical value ϵ . Obviously, if ϵ is set to zero, then there would not be any recurrence at all and, on the other hand, if ϵ is sufficiently large, then every trajectory would be considered a recurrence.

Hence, an *RP* is a graphical representation of the square matrix of the distances between all paired time coordinates of the reconstructed phase space, where a point in the distance matrix is darkened if the distance is smaller than ϵ and not otherwise. In other words, a Heaviside function is applied to the distances of each pair of time coordinates obtained by the reconstructed phase space. By construction, the main diagonal of the distances matrix is always darkened (and it is called *line of identity* - *LOI*), and

the idea of the *RP* is that if there is a significant amount of determinism in the system, then its phase-space trajectory will visit previous regions and, thus, the *RP* will show lines parallel to the main diagonal (*LOI*).

Marwan [11] summarizes three different structures that can be found in an *RP*, namely isolated dots, vertical lines, and diagonal lines. Single dots appear when states are rare or do not last enough time. A diagonal line occurs when a trajectory visits the same region of phase space that it was in the past. Finally, a vertical line appears when the state of the system does not move, or moves too slowly, for a period of time.

Figure 1 illustrates the *RP* of distinct time series patterns. The data were generated by mathematical models of a random, chaotic, periodic, and linear series, following a similar approach of Takakura et al. [17]. The random data was built using random integers between 0 and 100. The chaotic data was built by the logistic equation (1) with $x_0 = 0.2$ and $r = 3.7$, corresponding to a chaotic state. The values were also multiplied by 100, excluding the decimal place. The periodic series was built by repeating the numbers 0 to 90 with step 4, for a total of 180 numbers. The linear series was constructed by means of a time series comprised of prime numbers from 2 to 1,070.

2.3 Complexity Measures Based on the Recurrence Plot

Having seen that the *RP* is a powerful tool for analyzing dynamic systems, the *RQA* deals with measures which attempt to quantify the insights provided by the *RPs*. For instance, the periodic component of the Logistic Map is captured by long diagonal lines. Zbilut and Webber [23] developed some measures based on the density and the diagonal structures of the *RP*, and Marwan [11] extended some of these measures considering vertical structures. In an *RP*, diagonal lines suggest that the system orbits two regions of the phase space in parallel for some time, while vertical lines indicate that only one region of the phase space is being visited during that period.

REC - Recurrence Rate

It consists of summing all darkened dots in the *RP* and dividing them by the total number of dots. Mathematically,

$$REC(m, \tau, \epsilon) = \frac{1}{N^2} \sum_{i,j=1}^N R_{i,j}^{m,\epsilon} \quad (2)$$

where N is the length of the *RP*, and R is the binary distance matrix.

As it can be seen in the equation above, *REC* measures the density of the *RP*. Actually, the *REC* definition coincides with that from correlation dimension [8]. It is worth stressing that the *REC* measure can be adjusted by the radius ϵ defined in the previous section. Obviously, if one sets $\epsilon = 0$, then the *RP* will have no darkened points at all. On the other hand, if ϵ is set sufficiently large, then all points in the *RP* will be darkened.

One then should set ϵ as small as possible to expose a sufficient number of recurrence structures [12]. Therefore, there are many options to set ϵ , e.g. a percentage of the maximum phase space diameter [14], a value that should not exceed 10% of the mean or the maximum phase space [23], or setting ϵ so that the *REC* is approximately 1% [24]. The later option was adopted in this paper, but setting ϵ for all series targeting *REC* to 5%.

DET - Determinism

This measure is calculated by dividing the sum of all darkened points that belong to a diagonal structure by the total number of darkened points. For deciding whether a recurrence point belongs or not to a diagonal structure, one has to additionally define a minimal line length l_{min} . The idea is that a deterministic system produces longer diagonal lines than stochastic ones. Thus, we have

$$DET(m, \tau, \epsilon, l_{min}) = \frac{\sum_{l=l_{min}}^N l p^\epsilon(l)}{\sum_{i,j=1}^N R_{i,j}^{m,\epsilon}} \quad (3)$$

where $p^\epsilon(l)$ is the frequency distribution of diagonal line lengths. It should be stressed that if $l_{min} = 1$ (that is, every isolated recurrence point belongs to a line of length 1), then $DET = 1$.

RATIO - Ratio between *DET* and *REC*

The *RATIO* measure is simply the ratio of *DET* and *REC*, as shown below

$$RATIO(m, \tau, \epsilon, l_{min}) = \frac{DET(m, \tau, \epsilon, l_{min})}{REC(m, \tau, \epsilon)} \quad (4)$$

In the extreme case where $l_{min} = 1$ (and, thus, $DET = 1$), *RATIO* is equal to the reciprocal of *REC*. According to [21], in certain circumstances, this measure can be used to study phase transitions, because *REC* can decrease without a correspondent change in *DET*.

***L* - Average Diagonal Length**

This measure gives an impression about how much time two trajectories remain close to each other, and is simply calculated by obtaining the average value of the line lengths defined by l_{min} . Therefore, we have

$$L(m, \tau, \epsilon, l_{min}) = \frac{\sum_{l=l_{min}}^N lp^{\epsilon}(l)}{\sum_{l=l_{min}}^N p^{\epsilon}(l)} \quad (5)$$

In this sense, the more often the system shows long periods of recurrent orbits, the longer the average length of the diagonal lines.

***L_{max}* - Maximal Diagonal Length**

This measure is simply given by the maximal diagonal line length in the RP. According to Webber and Zbilut [21], it is related to the largest positive Lyapunov exponent.

$$L_{max}(m, \tau, \epsilon, l_{min}) = \max_{l=l_{min}}^N lp^{\epsilon}(l) \quad (6)$$

***DIV* - Divergence**

Divergence is the reciprocal of L_{max} . The smallest the divergence, the longer is the maximal recurrent orbit in the system.

$$DIV(m, \tau, \epsilon, l_{min}) = \frac{1}{L_{max}} \quad (7)$$

***ENT* - Entropy**

This measure consists of calculating Shannon entropy to the frequency distributions of diagonal line lengths and interpreting this as the complexity of the deterministic structure of the system. Thus, we have

$$ENT(m, \tau, \epsilon, l_{min}) = - \sum_{l=l_{min}}^N P(l) \ln P(l) \quad (8)$$

where

$$P(l) = \frac{p^{\epsilon}(l)}{\sum_{l=l_{min}}^N p^{\epsilon}(l)} \quad (9)$$

In this sense, ENT indicates the diversity of visited regions of the phase space. For instance, if the system visits very often a few specific regions of the phase space, and only occasionally lots of others regions, ENT will be low. In the extreme case of all regions being equally visited, ENT reaches its maximum value given by $N \log(N)$. However, as stated by Marwan [11], ENT depends heavily on l_{min} (that is, the bin size for creating the frequency distribution of line lengths).

***LAM* - Laminarity**

This measure is analogous to DET , but with regard to vertical lines. Vertical structures in a Recurrence Plot indicate stationarity in a specific region of the phase space. It is calculated by dividing the sum of all darkened points that belong to a vertical structure by the total number of darkened points. Thus, we have

$$LAM(m, \tau, \epsilon, v_{min}) = \frac{\sum_{v=v_{min}}^N vp^{\epsilon}(v)}{\sum_{i,j=1}^N R_{i,j}^{m,\epsilon}} \quad (10)$$

where v_{min} is the minimal length to define a vertical line.

***V_{max}* - Maximal Vertical Length**

This measure is simply given by the maximal vertical line length in the Recurrence Plot.

$$V_{max}(m, \tau, \epsilon, v_{min}) = \max_{v=v_{min}}^N vp^{\epsilon}(v) \quad (11)$$

More about those measures can be found in Marwan [11] or even in Trulla et al. [20] if related to the logistic map. In this paper, they are used in the context of surrogate non-linearity tests.

As an illustration, table 1 presents the selected RQA complexity measures for a random, chaotic, periodic, and linear time series presented in figure 1. And figure 2 shows the time series generated by the logistic map with four different values of the growth parameter, their respective RP , and selected RQA measures. The respective calculated values are shown in table 2. It can be seen that for the stable value $r = 3.52$ the majority of the RQA measures are different from the other three chaotic values.

Table 1: Selected RQA complexity measures for the random, chaotic, periodic and linear time series presented in figure 1

	<i>random</i>	<i>chaotic</i>	<i>periodic</i>	<i>linear</i>
<i>REC</i>	0.034	0.066	0.044	0.031
<i>DET</i>	0.051	0.801	1.0	0.909
<i>L_{max}</i>	3	18	157	98
<i>DIV</i>	0.333	0.056	0.006	0.01
<i>RATIO</i>	1.503	12.201	22.946	28.991
<i>LAM</i>	0.061	0.072	0.0	1.0
<i>V_{max}</i>	2	7	1	11
<i>ENT</i>	0.305	1.875	1.946	2.211

Table 2: Selected RQA complexity measures for four time series generated by different growth parameters of the logistic map. It can be seen that for the stable value $r = 3.52$ the majority of the RQA measures are different from the other three chaotic values.

	$r = 3.52$	$r = 3.679$	$r = 3.72$	$r = 4$
<i>REC</i>	0.224	0.072	0.051	0.039
<i>DET</i>	0.987	0.817	0.751	0.552
<i>L_{max}</i>	162	19	13	7
<i>DIV</i>	0.006	0.053	0.077	0.143
<i>RATIO</i>	4.413	11.281	14.807	14.285
<i>LAM</i>	0.0	0.536	0.092	0.005
<i>V_{max}</i>	1	8	9	3
<i>ENT</i>	2.568	2.783	2.762	1.786

2.4 Surrogate Non-Linearity Test Based on Recurrence Quantification

According to Marwan [11], a satisfying theoretical study on the properties of the measures described in the previous section is yet to be developed. However, with regard to stationary time series, significance levels of RQA measures can be assessed by means of surrogate tests. This tests are done in two steps.

First, samples of data generated and following a number of surrogates copies are produced for each sample. Here, the surrogate copies are generated by the method described in Schreiber and Schmitz [15] called *Iterative Amplitude Adjusted Fourier Transform* (IAAFT). The copies have the same autocorrelations and the same probability distribution as the source. The surrogate copies tends to preserve the linearity (overall form and short-lag relationships) and to destroy any nonlinearity in the original sequence. An excellent review can be found in Schreiber and Schmitz [16]. Second, the null hypothesis of linearity is rejected with regard to one of the RQA measures in each sample if all surrogated copies present values smaller than the measure in original data. As well pointed in Belaire-Franch [1], who also applied the same test on artificial financial data generated by Lux-Marcehshi model [9], this method prevents false rejections due to deviations from Gaussianity.

3 Data Series

The three data series submitted to a Recurrence Quantification Analysis (RQA) are presented below. The following subsection is dedicated to explaining the SSV microscopic model [6] of a financial market that is used as the generating process for the artificial data considered here. In the following, two real-world financial time series with very distinct behavior are presented.

3.1 Structural Stochastic Volatility Model (SSV)

With regard to agent design, this is a two-group model where agents can be fundamentalists or chartists. The main difference is that fundamentalists respond to deviations from fundamental price, and chartists extrapolate the returns they just observed in the previous period. Thus, their demand functions $d_t^{f,c}$ are given by

$$d_t^f = \phi(p^* - p_t) + \varepsilon_t^f \quad \varepsilon_t^f \sim N(0, \sigma_f^2), \quad \phi > 0 \quad (12a)$$

$$d_t^c = \chi(p_t - p_{t-1}) + \varepsilon_t^c \quad \varepsilon_t^c \sim N(0, \sigma_c^2), \quad \chi \geq 0 \quad (12b)$$

where the superscripts f and c denote agent affiliation (fundamentalists and chartists, respectively); the subscript t represents time unit; p is the log of the price; p^* is the log of the (fixed) fundamental price; $\varepsilon_t^{f,c}$ are group-specific random terms (with zero mean and $\sigma_{f,c}$ standard deviations) that account for intra-group heterogeneity; ϕ corresponds to the responsiveness of the fundamentalists to the deviation from fundamental price; and χ corresponds to the responsiveness of the chartists to the last trend observed.

However, this model also accounts for learning, in the sense that agents can dynamically change their minds and move to the other group. Therefore, the shares of each group in the total population are allowed to vary over time. Considering that $n_t^{f,c}$ denotes their respective population shares, total excess demand can be written as $n_t^f d_t^f + n_t^c d_t^c$. Seen that this equation may not balance, a market maker is assumed to hold a sufficiently large inventory for supplying any excess of demand and for absorbing

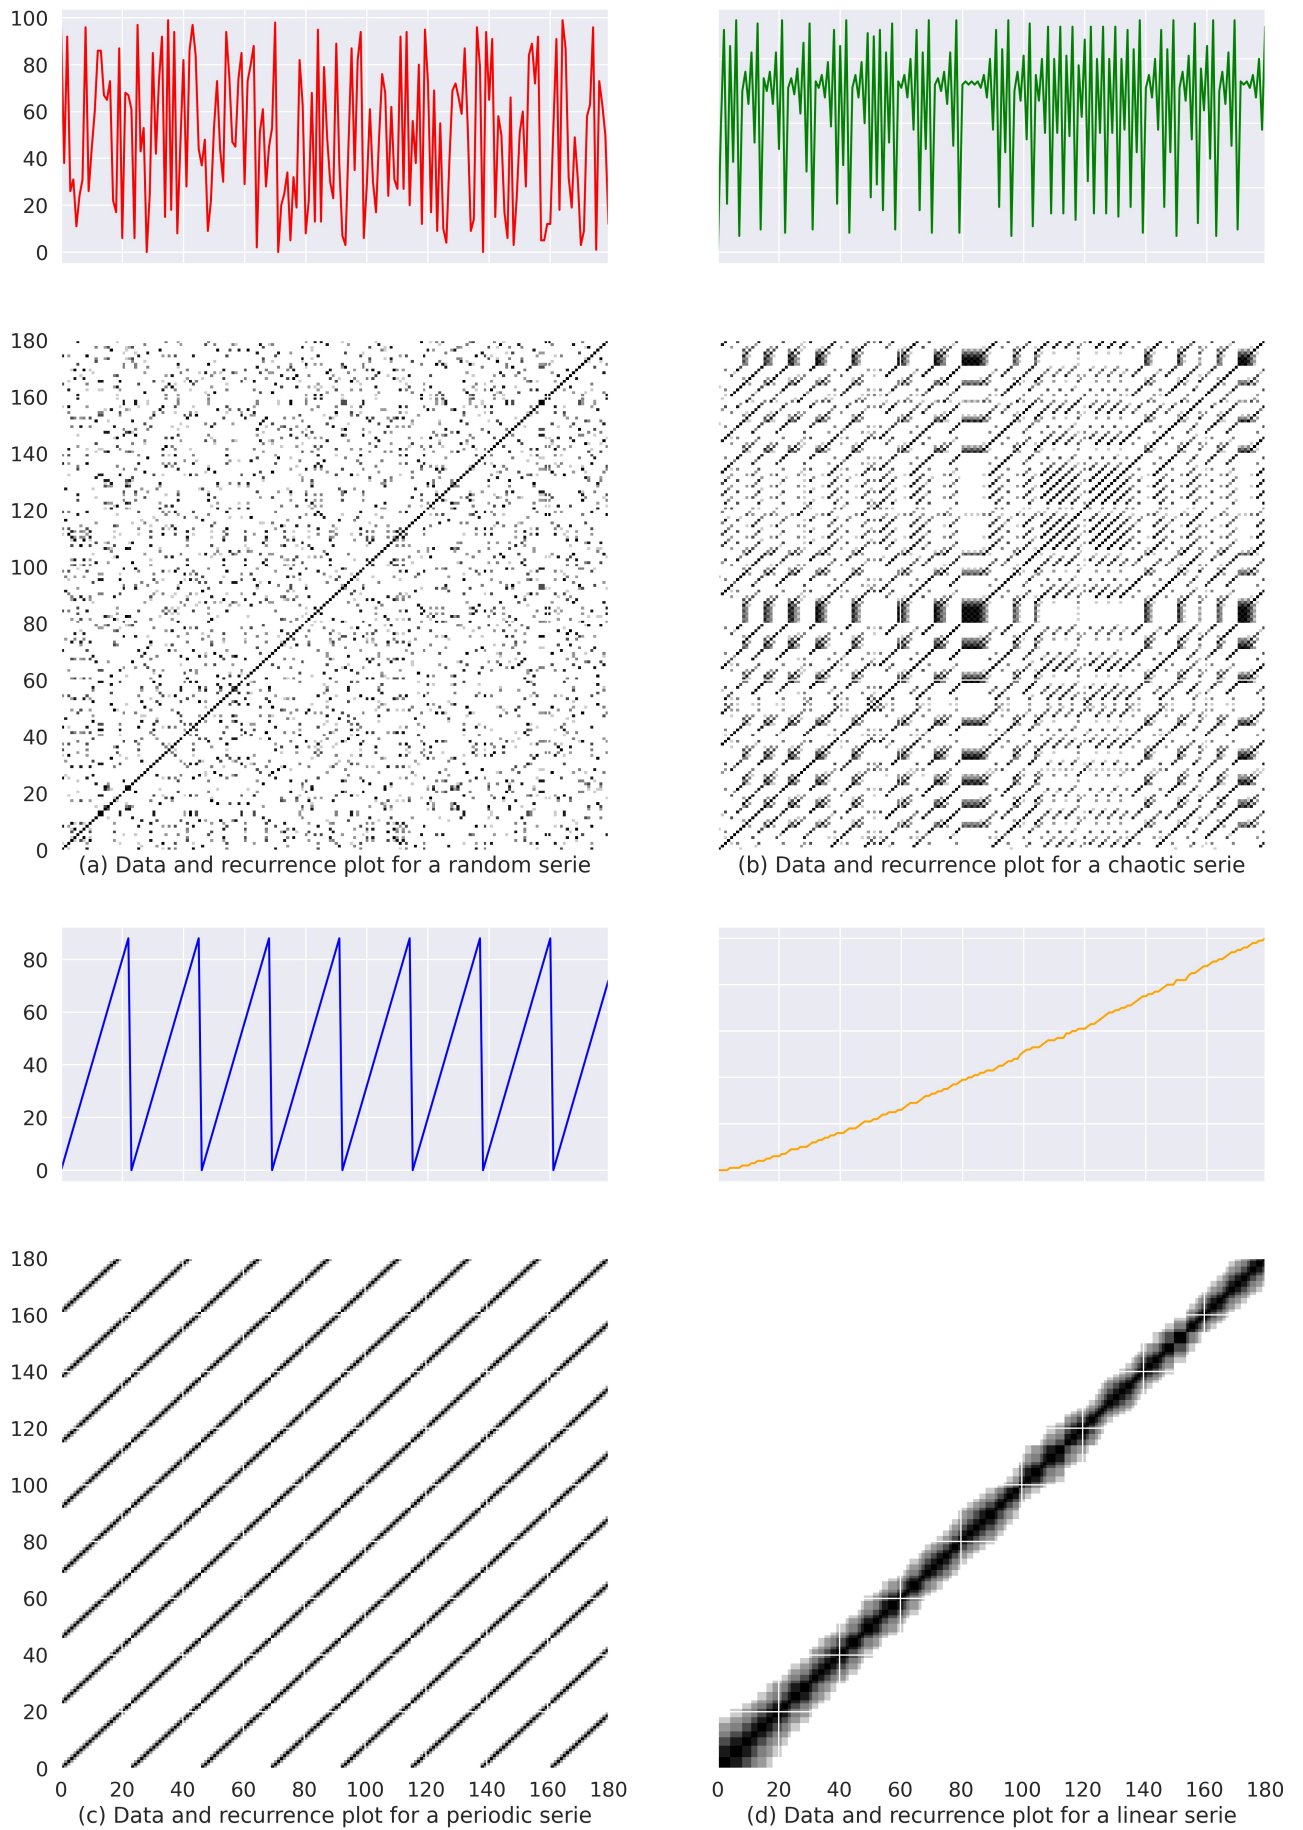


Figure 1: Time series generated by four different patterns and their respective recurrence plots.

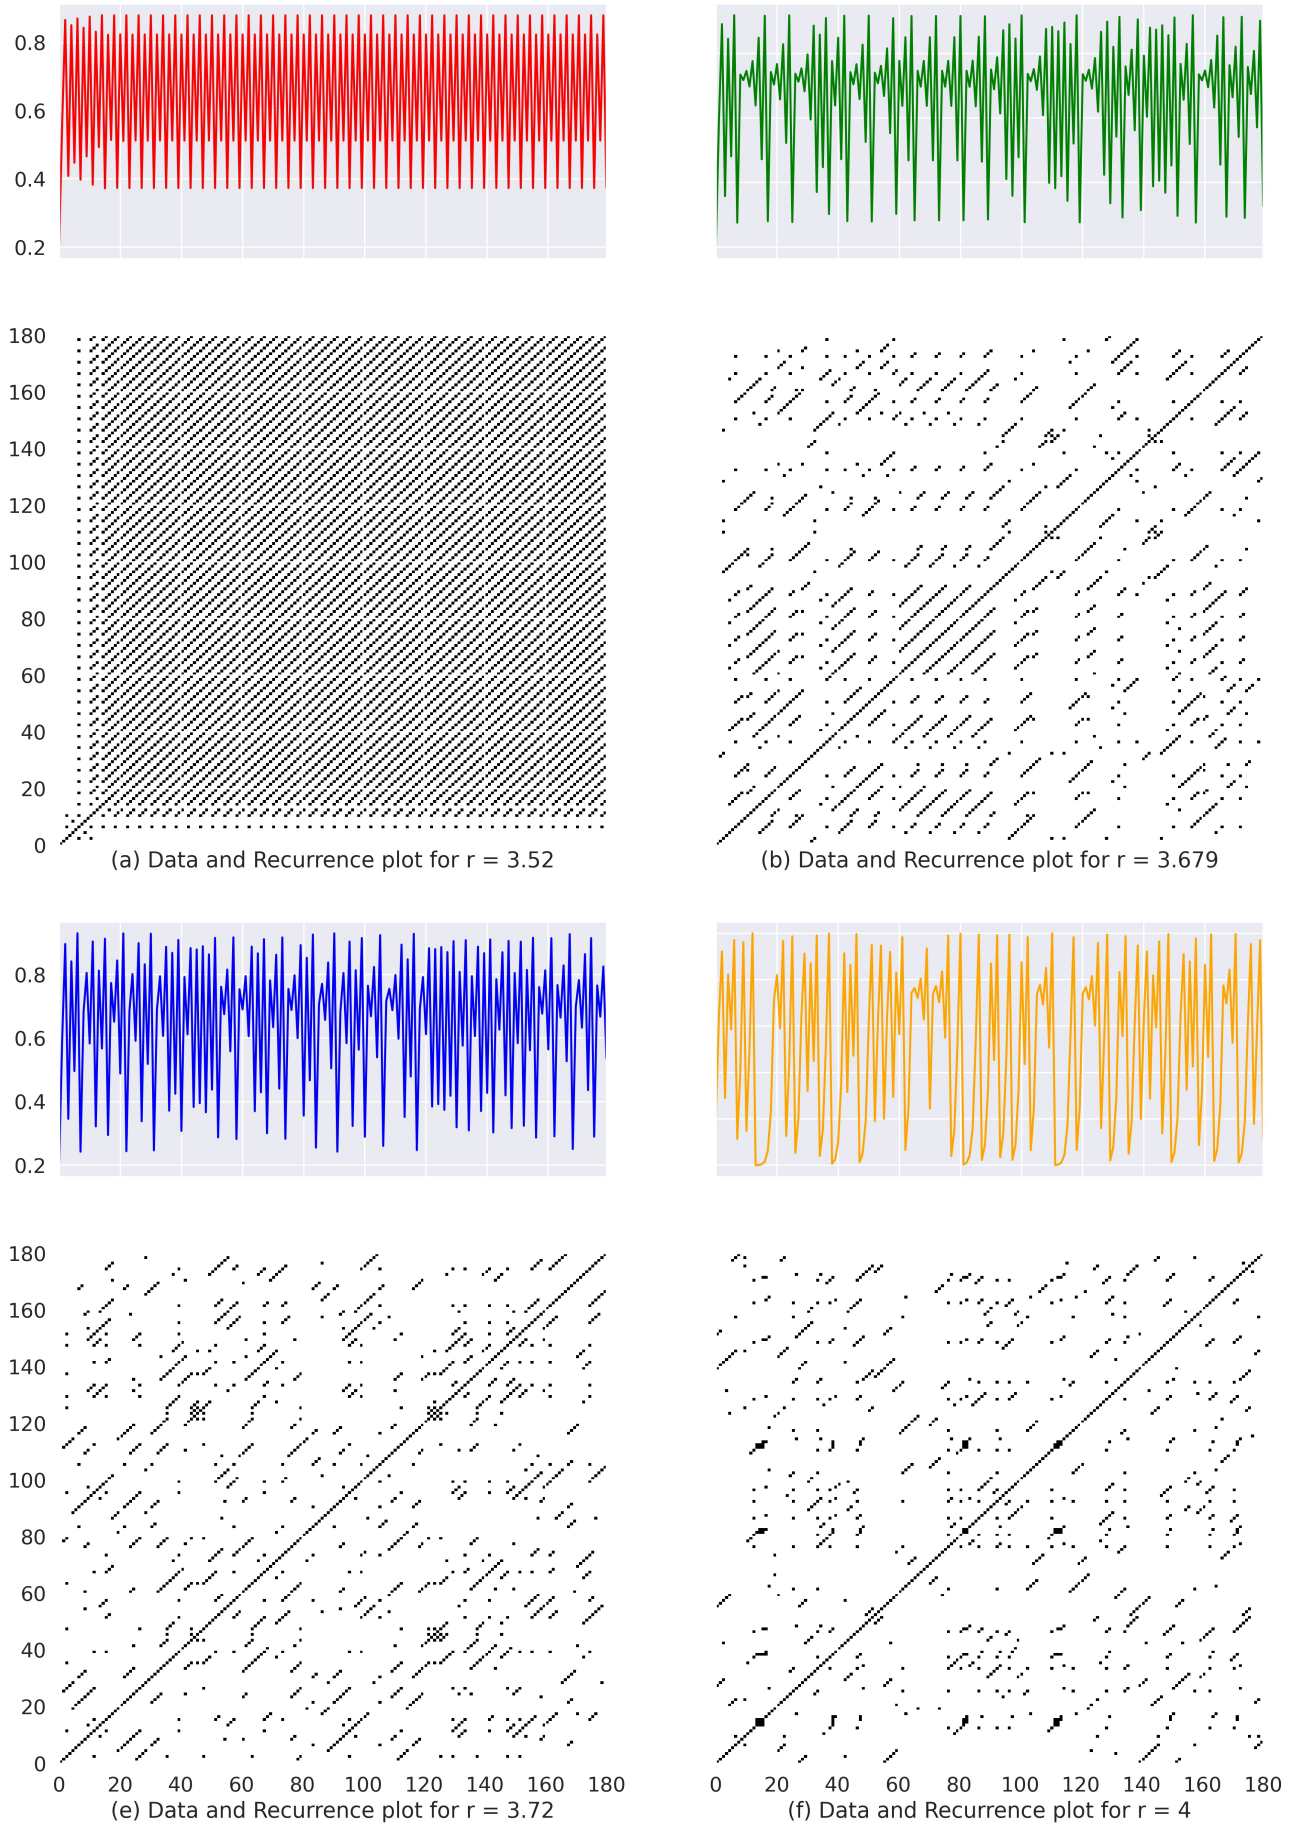


Figure 2: Time series generated by four different growth parameters of the logistic map and their respective recurrence plots.

any excess of supply. This is done by adjusting the price in the next period by a fixed coefficient that is inversely related to market liquidity. Considering these specifications, price determination at each time unit is given by

$$P_{t+1} = P_t + \mu[n_t^f \phi(p^* - p_t) + n_t^c \chi(p_t - p_{t-1}) + \varepsilon_t] \quad (13)$$

where

$$\varepsilon_t \sim N(0, \sigma_t^2), \quad \sigma_t^2 = (n_t^f)^2 \sigma_c^2 + (n_t^c)^2 \sigma_c^2 \quad (14)$$

summarizes what the authors coined as *Stochastic Structural Volatility* (SSV), and can be viewed as a structural modeling approach to time-varying variance.

What remains to be explained is the learning mechanism that yields the dynamics of the population shares. Even though the authors presented two different technical approaches for this, namely transition probabilities and discrete choice, only the latter will be considered here, given its best performance in a comparative study conducted by the same authors [6]. It is worth beginning with the definition of a switching index s_t , which attempts to measure the relative attractiveness of the fundamentalist's strategy in comparison to that of the chartist, given by

$$s_t = \alpha_0 + \alpha_x x_t + \alpha_d (p_t - p^*)^2 \quad (15)$$

where α_0 is a predisposition parameter to switch to fundamentalism; $\alpha_x x_t$ captures the idea of herding behavior; x_t is equal to $n_t^f - n_t^c$ in order to capture the relative presence of the two groups; and α_d can be understood as a measure of the influence of price misalignment (that is, the larger the gap, the higher the attractiveness of switching to fundamentalism). Thus, in the spirit of the discrete-choice approach, the shares of the total population in each group can be written as $n_t^f = 1/[1 + e^{-\beta s_{t-1}}]$ and $n_t^c = 1 - n_t^f$, where β is the intensity of choice. Figure 3 compares outputs from a single run of the model with returns of Standard & Poor's 500 index as a benchmark.

3.2 Real-world benchmarks

As one of the purposes of this paper is to contrast artificial and real-world *RQA* applications, two very contrasting real-world scenarios were selected for the sake of generalization. In addition to the artificially generated data from the *SSV* model described above, the *RQA* methodology is applied to daily returns from the Standard & Poor's 500 index (closing values) and from the market price in dollars of one Bitcoin (latest data in UTC time). The three time-series were limited to 2,000 observations, and are depicted both in level and in percent returns in figure 4.

As it is clear from the figure, the market value of one Bitcoin has experienced extreme variation during the time frame considered here. Such a fast and expressive value increase imprints on the time series a consistent upward trend during almost half of the time period considered. On the other hand, the S&P index presents alternating periods of ups and downs of approximated magnitude, with the exception of the high volatility period at the center.

4 Surrogate Test Results

This section presents details of the application of surrogated non-linearity tests on the data generated by the *SSV* model [6] and on the real-world data series. The first subsection reproduces the approach used in Belaire-Franch [1], while the second uses a slightly different approach in order to compare *RQA* of real-world and artificial financial data.

4.1 Testing for non-linear structures in artificial financial data

For the sake of comparability, the same approach in Belaire-Franch [1] was followed. This means that a total sample size of 40,000 observations were generated by the *SSV* model, and divide into 20 subsamples of size 2,000 each. For each of these subsamples, 20 surrogate copies were produced, and four of the *RQA* measures, namely *REC*, *DET*, *ENT* and L_{max} , were then compared with the value from the original subsamples. The parameters for those tests are $m = 15$, $\tau = 1$, $l_{min} = 3$ and ϵ is endogenously determined to aim a fixed 5% REC in the original series. The null hypothesis of linearity is rejected with regard to one of these measures if all surrogated copies from a specific subsample present values smaller than its respective subsample. Results are summarized in table 3.

As it can be seen in table 3, results from different test run are not similar within subsamples. Interestingly, while comparing test results with the visual appearance of the relevant parts of the time series (Fig. 5), there seems to be a general tendency towards rejection in periods with larger fluctuations, while in periods with moderate volatility linearity is not rejected.

For instance, in the same way as pointed out by Chen et al. [2] and Belaire-Franch [1], both acceptance and rejection of linearity are found in different realizations of the same model. However, different complexity measures usually agree within subsamples, suggesting both that the model is able to reproduce the empirically observed alternating periods of high and low volatility and also that the *RQA* measures are able to discriminate between these subsamples, as shown in figure 5. It can be seen that, for the *SSV* model, the null hypothesis of linearity is more likely to be accepted if subsample volatility is low, and rejected otherwise.




Figure 3: $T = 6750$ observations of (a) log of price, (b) share of fundamentalists, and (c) returns from a simple run of the model and (d) daily returns from S&P500 from January 1980 to March 2007. Inputs to the model are as follows: $\phi = 0.0728$, $\chi = 0.0896$, $\mu = 0.01$, $\alpha_0 = -0.327$, $\alpha_x = 1.815$, $\alpha_d = 9.6511$, $\sigma_f = 1.0557$, $\sigma_c = 2.9526$, $p^* = 0$, and $\beta = 1$.




Figure 4: Artificially generated data from the *SSV* model (a), daily returns from the Standard & Poor's 500 index (b) and from the market price in dollars of one Bitcoin (c). The three time-series were limited to 2,000 observations, and are depicted both in level and in percent returns

Table 3: Results of non-linearity tests based on RQA complexity measures for the SSV model. It can be seen that results from different tests are similar within most sub-periods. Comparing test results with the visual appearance of the relevant parts of the time series (Fig. 5), there seems to be a general tendency towards rejection in periods with larger fluctuations.

Subsample	REC	DET	ENT	L_{max}
1	Reject	Reject	Reject	Reject
2	Reject	Accept	Reject	Accept
3	Reject	Reject	Reject	Reject
4	Reject	Reject	Reject	Reject
5	Reject	Reject	Reject	Reject
6	Reject	Reject	Reject	Reject
7	Reject	Reject	Reject	Reject
8	Reject	Reject	Reject	Accept
9	Reject	Reject	Reject	Reject
10	Reject	Reject	Reject	Reject
11	Reject	Reject	Reject	Reject
12	Reject	Reject	Reject	Reject
13	Reject	Reject	Reject	Reject
14	Reject	Reject	Reject	Reject
15	Reject	Reject	Reject	Reject
16	Reject	Accept	Reject	Accept
17	Reject	Reject	Reject	Reject
18	Reject	Reject	Reject	Reject
19	Reject	Reject	Reject	Accept
20	Accept	Accept	Accept	Accept




Figure 5: Artificial returns for the SSV model. It can be seen in table 3 that results from different tests are similar within most subperiods. Comparing test results with the visual appearance of the relevant parts of the time series, there seems to be a general tendency towards rejection in periods with larger fluctuations.

Figures 6 and 7 show Recurrence Plots of two subsamples from the time series depicted in figure 5. For instance, considering the *SSV* model, the null hypothesis of linearity is rejected by all measures in the subsample 5, and not rejected by all measures in subsample 20. It can be seen that, as also found by Chen et al. [2] and Belaire-Franch [1], the null hypothesis of linearity is often rejected in periods presenting high volatility dominated by speculative trading. During periods of low volatility, lower complexity is captured by the *RQA* measures and, hence, it is easier to accept the null hypothesis of linearity.

In addition, it should be pointed out that the hypothesis of chaos or linearity is rejected for the majority of the subsamples. This indicates that, if there is a deterministic process ruling the data, it is more complicated than the dynamics from low dimensional chaotic systems. When considering the *RQA* results for the Lux-Marchesi model presented by Belaire-Franch [1], approximately the same number of rejection is observed here, although a few more occur for the *SSV* model when compared to the Lux-Marchesi model. This is somewhat expected since the *SSV* family of models is known for better reproducing financial time series [6]. For instance, the null hypothesis of linearity is rejected here 19 out of 20 with respect to the complexity measure *REC* (recurrence), while in the results presented by the aforementioned authors this number is 18. The number of rejections found here (17) is above theirs (15) when considering the *DET* (determinism) measure. When considering the *ENT* (entropy) measure the results presented here show 2 rejections less than the authors (19 against 17), and when considering the L_{max} (maximum diagonal length) measure this difference is of 2 rejections (15 and 13).

4.2 Testing for non-linear structures in real-world financial data

In this second experiment, some modifications were necessary in order to compare such results to those obtained from applying *RQA* to real-world financial data. Particularly, the sample size was restricted to 2,000 (approximately 5.5 years of daily data), no division in sub-samples was performed, but 100 surrogated copies were generated instead.

This time the null hypothesis is rejected only if the value of the measure from the original series is significantly different from the ones generated from the surrogate set. This approach is based on the idea that surrogate data testing is a statistical proof by contradiction technique similar to parametric bootstrapping [19]. In this sense, the proportion of surrogate copies presenting an *RQA* measure higher than its correspondent series is interpreted as a p-value for hypothesis testing. Once more, to make all the *RP* complexity measures comparable among the series, the radius ϵ was set to guarantee that *REC* of all series and copies to be 5%.

Table 4: The proportions of surrogate copies presenting an *RQA* measure higher than its correspondent series.

Metric	Bitcoin	S&P500	SSV
<i>DET</i>	0.0	0.00	0.00
<i>DIV</i>	1.0	1.00	0.78
<i>ENT</i>	0.5	0.01	0.03
<i>LAM</i>	0.0	0.00	0.75
L_{max}	0.0	0.00	0.24
V_{max}	0.0	0.00	0.34

The range of *RQA* measures applied in this experiment was extended, to include *DIV*, *LAM*, and V_{max} in addition to the previous *REC*, *DET*, *ENT*, and L_{max} measures. Of course, *REC* is not considered because all series and copies have it fixed at the 5% level.

The parameters for generating the *RPs* were kept at the previous values of $m = 15$, $\tau = 1$, $l_{min} = 3$, and ϵ being endogenously determined to target a fixed 5% *REC*. However, here such procedure was applied both for the original and the surrogate copies time series. The proportion of surrogate copies presenting values above its corresponding original series is presented in table 4 for each *RQA* measure.

It can be seen that the null hypothesis of linearity has been rejected for all three series only by the *DET* and *RATIO* measures. With respect to *DET*, it was not able to detect any determinism in data generated by the *SSV* model neither in the real-world financial series.

However, for the same time series, the proportions of copies with values above the original for the L_{max} , *LAM*, and V_{max} do not allow the null hypothesis of linearity to be rejected (at least not under a reasonable significance level for the *SSV* data), meaning that some structure was detected in the data generated by the model, and it was not present in the real-world time series. The *SSV* is a more evolved model when compared with the one used in Belaire-Franch [1], the Lux-Marchesi model [9], but it presents a deterministic core.

In this sense, those measures were able to tell the artificial and the real-world series apart, even being the SP&500 and the Bitcoin indexes so different in statistical terms. Noticing that *DIV* is the reciprocal of L_{max} , the null hypothesis could not be rejected for the three series by both measures, neither helped to discriminate between artificial and real-world series. With respect to the *ENT* measure, its p-value reveals an interesting pattern by capturing the fact that the Bitcoin has a steady uptrend, and it was enough to avoid its null hypothesis rejection for linearity.

5 Conclusion

The main interest in this work is twofold. First, to extend the results presented by Chen et al. [2] and Belaire-Franch [1] using the Recurrence Quantification Analysis (*RQA*) in data generated by a more evolved microscopic model (*SSV*). Second, to show how to apply and compare *RQA* complexity measures among different data, including artificial and real-world data. More specifically, artificial data generated by the *SSV* model and two financial data series with very distinct behavior by nature: Bitcoin, and S&P500.

In the second section, the *RQA* approach was shown to be powerful by allowing one to discriminate between the different periods of periodic orbits and chaotic motion of the logistic map.

Following the introduction of the Structural Stochastic Volatility (*SSV*), the results show that the *RQA* approach combined with a surrogate linearity test is able to discriminate between periods of high and low volatility in artificial financial data, in the same way as pointed out by Chen et al. [2] using traditional non-linearity tests, and later by Belaire-Franch [1] addressing data from the Lux-Marchesi model. In addition, an approach for making *RQA* measures comparable among distinct series was explored, revealing that *RQA* measures are also a valuable resource that can be used to characterize artificial and real-world series.

Again in agreement with Chen et al. [2], it can be seen that the hypothesis of linearity is rejected for the majority of the subsamples of data generated from the Structural Stochastic Volatility (*SSV*) model. In this sense, this work extends and supports the results presented by Chen et al. [2] using a different non-linearity test based on the *RQA*, and by checking if such results are confirmed in a different artificial market model (*SSV*). Finally, the results add evidence in agreement with Chen et al. [2] and Belaire-Franch [1] to support the rejection of linearity or low dimensional chaotic motion in an artificial financial time series generated from the *SSV* microscopic model, which is a model that better replicate dysfunctions in financial markets.

To conclude, previous work tested for non-linearity or the presence of low dimensional chaos both in artificial and real-world financial time series. The contribution of this paper to this research is twofold: first to contrast previous findings with the application of the *RQA* methodology with data generated by a more evolved microscopic model of artificial market, and second to evaluate the *RQA* results obtained above with the same methodology being applied to real-world financial data, in order to assess if *RQA* results could help to distinguish between artificial and real-world data, even if linearity is rejected in both cases. It is shown evidence, in agreement previous findings, to support the rejection of linearity or low dimensional chaotic motion in an artificial financial time series generated from the *SSV* microscopic model. In addition, it is also shown that that *RQA* measures can help to discriminate artificial from real-world financial data, at least when specific *RQA* measures are considered.




Figure 6: Subsample number 5 and its Recurrence Plot from the *SSV* model time series depicted in figure 5. The difference between this subsample and subsample number 20 is highlighted in their Recurrence Plots and in table 3. In subsample number 5, the null hypothesis of linearity is rejected by all four measures, while in the subsample number 20 the null is not rejected by any of the measures. It can be seen that, as also found by [1] and [2], the null hypothesis of linearity is more often rejected in periods presenting high volatility dominated by speculative trading. During periods of low volatility, the lower complexity of the time series is captured by the *RQA* measures, thus, making it is easier to the null hypothesis of linearity for being accepted.




Figure 7: Subsample number 20 and its Recurrence Plot from the *SSV* model time series depicted in figure 5. The difference between this subsample and subsample number 5 is highlighted in their Recurrence Plots and in table 3. In subsample number 20, the null hypothesis of linearity is not rejected by any of the four measures, while in the subsample number 5 the null is rejected by all the measures. It can be seen that, as also found by [1] and [2], the null hypothesis of linearity is more often rejected in periods presenting high volatility dominated by speculative trading. During periods of low volatility, the lower complexity of the time series is captured by the *RQA* measures, thus, making it is easier to the null hypothesis of linearity for being accepted.

References

- [1] J. Belaire-Franch. Testing for non-linearity in an artificial financial market: a recurrence quantification approach. *Journal of Economic Behavior and Organization*, 54:483–494, 2004.
- [2] S. Chen, T. Lux, and M. Marchesi. Testing for non-linear structure in an artificial financial market. *Journal of Economic Behavior and Organization*, 46(3):327–342, 2001.
- [3] F. Cont. Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1(1):223–236, 2001.
- [4] J. Eckmann, S. Kamphorst, and D. Ruelle. Recurrence plots of dynamical systems. *Europhysics Letters*, 4(9):973–977, 1987.
- [5] A. Fabretti and M. Ausloos. Recurrence plot and recurrence quantification analysis techniques for detecting a critical regime. *International Journal of Modern Physics C*, 16(5):671–706, 2005.
- [6] R. Franke and F. Westerhoff. Structural stochastic volatility in asset pricing dynamics: Estimation and model contest. *Journal of Economic Dynamics and Control*, 36(8):1193–1211, 2012.
- [7] R. Giglio. *Essays on Interlocking Directorates and Speculative Dynamics*. PhD thesis, Kiel University, 2015. URL https://macau.uni-kiel.de/receive/diss_mods_00017587.
- [8] P. Grassberger and I. Procaccia. Characterization of strange attractors. *Phys. Rev. Lett.*, 50:346–349, 1983.
- [9] T. Lux and M. Marchesi. Scaling and criticality in a stochastic multi-agent model of a financial market. *Nature*, 397(6719):498, 1999.
- [10] B. G. Makiel. Efficient market hypothesis. In N. P. Eatwell J., Milgate M., editor, *Finance*, pages 127–134. The New Palgrave, Palgrave Macmillan, London, 1989.
- [11] N. Marwan. *Encounters With Neighbours – Current Developments Of Concepts Based On Recurrence Plots And Their Applications*. PhD thesis, University of Potsdam, 2003.
- [12] N. Marwan. How to avoid potential pitfalls in recurrence plot based data analysis. *International Journal of Bifurcation and Chaos*, 21(4):1003–1017, 2011.
- [13] R. May. Simple mathematical models with very complicated dynamics. *Nature*, 261(5560):459–467, 1976.
- [14] G. Mindlin and R. Gilmore. Topological analysis and synthesis of chaotic time series. *Physica D: Nonlinear Phenomena*, 58(1):229–242, 1992.
- [15] T. Schreiber and A. Schmitz. Improved surrogate data for nonlinearity tests. *Physical Review Letters*, 77:635–638, 1996.
- [16] T. Schreiber and A. Schmitz. Surrogate time series. *Physica D: Nonlinear Phenomena*, 142(3):346–382, 2000.
- [17] I. T. Takakura, R. A. Hoshi, M. A. Santos, F. C. Pivatelli, J. H. Nóbrega, D. L. Guedes, V. F. Nogueira, T. Q. Frota, G. C. Castelo, and M. F. de Godoy. Recurrence plots: a new tool for quantification of cardiac autonomic nervous system recovery after transplant. *Brazilian Journal of Cardiovascular Surgery*, 32(4):245–252, 2017.
- [18] F. Takens. *Detecting strange attractors in turbulence*. In: *Dynamical systems and turbulence. Warwick 1980*. Springer, Berlin, Heidelberg, 1981.
- [19] J. Theiler, S. Eubank, A. Longtin, B. Galdrikian, and J. Farmer. Testing for nonlinearity in time series: The method of surrogate data. *Physica D: Nonlinear Phenomena*, 58:77–94, 1992.
- [20] L. Trulla, A. Giuliani, J. Zbilut, and C. Webber. Recurrence quantification analysis of the logistic equation with transients. *Physics Letters A*, 223(4):255–260, 1996.
- [21] C. Webber, Jr and J. Zbilut. Dynamical assessment of physiological systems and states using recurrence plots strategies. *Journal of Applied Physiology*, 76(2):965–973, 1994.
- [22] J. Zbilut and C. Webber. Recurrence quantification analysis: Introduction and historical context. *International Journal of Bifurcation and Chaos*, 17(10):3477–3481, 2007.
- [23] P. Zbilut and C. Webber, Jr. Embeddings and delays as derived from quantification of recurrence plots. *Physics Letters A*, 171(3-4):199–203, 1992.
- [24] P. Zbilut, J. Zaldivar-Comenges, and F. Strozzi. Recurrence quantification based liapunov exponents for monitoring divergence in experimental data. *Physics Letters A*, 297(3-4):173–181, 2002.