Time series clustering based on time-varying Hurst exponent

Alex Babiš*, Beáta Stehlíková

Comenius University, Faculty of Mathematics, Physics and Informatics, Bratislava, Slovakia

Abstract

We consider the problem of clustering time series which are assumed to possess the long term memory. We propose an approach based on combining the results obtained by applying different methods for estimating time-varying Hurst exponent and apply it to Euro exchange rates. Firstly, we fit AR-GARCH models to every time series to reduce bias of rescaled range analysis method. We only consider model with residuals, in which no autocorrelation and ARCH effect is present; among them we choose the model with the lowest value of the Bayesian information criterion. Afterwards, we estimate the Hurst exponent from the residuals by means of the rolling window approach using four different estimation methods. Vectors of Hurst exponents are clustered for each of the four cases and the clusters are compared in order to obtain the final clustering.

Keywords: Hurst exponent, Clustering, Stock market, Time series, GARCH

1. Introduction

Clustering, also known as cluster analysis, is an important tool for analysing data. Clustering methods partition data into several homogeneous groups called clusters. Clusters are created so that similarity between object within specific cluster is maximized, while at the same time similarity between objects that do not belong to the same cluster is minimized. Clustering can be used as a part of exploratory data analysis, as it allows us to gain information from underlying data without explicit knowledge about relationship between objects within. It can also provide useful insights on the structure of the data and identification of groups containing similar observations might be of interest on its own. Usually, information about objects is given by a vector of features. In many applications a vector of features arises by observing specific characteristics of an object at different time intervals. Resulting vectors have therefore the form of time series.

Clustering of time series data found its way into a wide range of areas, such as astronomy, medicine, environmental analyses, etc. We refer the reader to review paper by Aghabozorgi et al. (2015) for more applications and concrete references. In finance, where also our dataset belongs, applications include for example clustering stocks with use in portfolio optimization

---

*Corresponding author

Email addresses: alexbabis96@gmail.com (Alex Babiš), stehlikova@fmph.uniba.sk (Beáta Stehlíková)

(Han & Ge, 2020; Iorio et al., 2018; Massahi et al., 2020) or clustering aiming to discover the structure of cryptocurrencies market (Song et al., 2019).

Similarly to general cluster analysis, there is a great number of different methods and approaches. Time series clustering can be based on clustering the data directly, on extracting their features or on model which were fit to the data. There are new distance metrics, proposed specifically for time series. More details can be found for example in survey papers (Aghabozorgi et al., 2015; Fu, 2011) or in a recent book by Maharaj et al. (2019).

Our approach is based on estimating the Hurst exponent of the time series, which is a measure of long-range dependence in the time series. Origin of the Hurst exponent dates back to 1951, when British hydrologist Harold E. Hurst proposed a method to optimize the storage capacity of reservoirs in an effort to regulate natural contribution of the Nile river, keeping in mind cyclical trend such as periods of drought and floods. His statistical analysis of the hydrological data was not in accord with standard models of that time and subsequently lead to models describing the behaviour that could be characterized as a long-range dependence or long memory (O'Connell et al., 2016). Applications of the Hurst exponent in finance include analyses of interest rates (Cajueiro & Tabak, 2009), hedge funds performance (Auer, 2016), energy futures market (Sensoy & Hacihasanoglu, 2014), cryptocurrencies (Jiang et al., 2018), efficiency of stock market (Cajueiro & Tabak, 2004) and others.

In the same way as Cajueiro and Tabak (2004), we apply the Hurst exponent estimators to standardized residuals from AR-GARCH model. It was shown that presence of short memory could cause bias of estimated value of Hurst exponent and using this procedure we filter out the short memory information. In this way our analysis of long-range memory is not affected by short memory effects present in the data. Paper by Lahmiri (2016) used different Hurst exponent estimators to cluster industrial sectors at Casablanca Stock Exchange. We follow a similar idea in our approach. However, instead of using a single estimate of the Hurst exponent for the whole time series, we use the rolling window approach. In other settings it has been successfully used, among others, in Cajueiro and Tabak (2004), Jiang et al. (2018), and Sensoy and Hacihasanoglu (2014). We use this approach for subsequently clustering the time series of Hurst exponent estimates.

Our contribution therefore lies in combining several approaches used in the literature dealing with financial time series and their Hurst exponents individually, but not in this combination—using residuals from AR-GARCH models for the estimation of the Hurst exponent, using rolling window estimates and clustering the time series. Furthermore, we propose a network based method for clustering, based on the results from an arbitrary number of clustering algorithms applied to the data. It can be used in a more general setting, not only our choices of methods for estimating the Hurst exponents and the clustering procedures.

The rest of the paper is organized as follows. In Section 2 we review the notion of the Hurst exponent and its estimators, which we will use in our analysis. Section 3 presents our data set and Section 4 summarizes the results of GARCH modelling applied the data. Section 5 shows the Hurst exponent estimates and their clusterings. In Section 6 we compare these clusters and suggest the final clustering of the time series. We conclude the paper with remarks on the methodology, its advantages and shortcomings, and with ideas for future research in Section 7.

2. Long-range dependence in time-series and estimation of the Hurst exponent

If the dependence between observations of a stationary time series that are far apart from each other decreases very slowly, as the time distance between them increases, then the

time-series is said to exhibit long-range dependence or long memory. More specifically, the autocorrelations $\rho(s) = Cor(X_t, X_{t-s})$ decay to zero so slowly, that they are not absolutely summable, i.e. $\sum_{s=0}^{\infty} |\rho(s)| = \infty$. This holds in contrast to ARMA models, for which the autocorrelation function decays exponentially and therefore the sum of its absolute values is finite. Typical long memory process have $\rho(s) \sim |s|^{-\alpha}$ with $\alpha \in (0,1)$, as $s \to \infty$. Other models of long memory processes include ARFIMA models with fractional differences (in contrast to integer differences in ARIMA models), they can be characterized via spectrum, or so called Hurst exponent. A detailed treatment of the long memory processes can be found in Beran (2017).

Hurst exponent, which we use in our analysis, attains the values from the interval $(0,1)$ and, if different from $1/2$, it is linked to the asymptotic behaviour of the autocorrelation function by the relation $\rho(s) \sim H(2H-1)|s|^{2H-2}$. The case $H = 1/2$ corresponds to processes with exponentially decaying autocorrelations, i.e. without the long memory. Values $H \in (1/2,1)$ correspond to persistent processes, while values $H \in (0,1/2)$ correspond to anti-persistent processes.

The oldest and probably the best-known method for estimation of the Hurst exponent is Rescaled range (R/S) analysis, originally proposed by Hurst (1951) himself and further developed by Mandelbrot and Wallis (1969). We outline this method according to Weron (2002) and afterwards we explain its modifications which we have used in our analysis, using their implementation in the R package `pracma` (Borchers, 2019), in particular the function `hurstexp()`.

Let $\{X_t\}_{t=1}^L$ be stationary time series of length L. The Hurst exponent can be estimated as follows:

1. Time series of length L is divided into d sub-series of length n.
2. For each sub-series, indexed by m, mean E_m and standard deviation S_m are calculated.
3. Data $X_{i,m}$ are than normalized by subtracting mean E_m:

$$\hat{X}_{i,m} = X_{i,m} - E_m \quad (i = 1, \ldots, n).$$

4. Next step is to calculate new time series of deviations from mean value for each sub-period:

$$Y_{i,m} = \sum_{j=1}^{i} \hat{X}_{j,m} \quad (i = 1, \ldots, n).$$

5. The range R_m is calculated as

$$R_m = \left\{ \max\{Y_{1,m}, \ldots, Y_{n,m}\} - \min\{Y_{1,m}, \ldots, Y_{n,m}\} \right\}.$$

6. Each range R_m is then rescaled/normalized by standard deviation for corresponding sub-period as $\frac{R_m}{S_m}$.
7. Finally mean value of the rescaled range for all sub-series of length n is computed

$$\frac{R}{S}(n) = \frac{1}{d} \sum_{m=1}^{d} \frac{R_m}{S_m}$$

8. The steps above are repeated for the increasing length n. Only the values of n which include first and last points of time-series are used, so $\frac{R}{S}(n)$ is calculated from the same number of observations for each n.

It was shown, (cf. Di Matteo, 2007; Mandelbrot, 1975; Mandelbrot & Wallis, 1969; Taqqu et al., 1995), that $\frac{R}{S}$ statistics asymptotically follows relation

$$\frac{R}{S}(n) \sim cn^H.$$

Taking logarithm leads to

$$\log\left(\frac{R}{S}(n)\right) \sim H\log(n) + \log(c). \tag{2.1}$$

It means that in order to estimate value of Hurst exponent H it is sufficient to run simple linear regression over sample of increasing time interval n.

The algorithm above has been modified in several ways in the literature. The simplest form of the rescaled range analysis would be not to separate original time-series into m sub-series but rather considered whole time series as suggested originally in Hurst (1951). This would lead to only one $\frac{R}{S}(n)$ statistics which means taking

$$\frac{\log\left(\frac{R}{S}(n)\right)}{\log(n)}$$

would be sufficient enough to estimate Hurst exponent H. This method is referred to as *simplified rescaled range analysis*.

Results of the rescaled range analysis can depend on the choice of the lengths of sub-series n that are used as input into regression. If the starting value of n is chosen as the length of the original time series and then progressively halved, then it would possibly mean, if $n \neq 2^i$ for some i, that last sub-series would be of different length as the all previous. The resulting statistics will be referred to as *corrected rescaled range analysis*.

A better way to estimate Hurst exponent H via classical rescaled range analysis would be to only consider those lengths of sub-series n that the length of the original series is multiple of n. This statistics will be referred to as *empirical rescaled range analysis*.

As stated in Annis and Lloyd (1976) and Peters (1994), for small value of n, the deviance of the slope in the regression (Equation (2.1)) from its true value is significant even in a simple case, when the underlying process is a Gaussian noise. They approximate the theoretical values for $\frac{R}{S}(n)$ as

$$E\left(\frac{R}{S}(n)\right) = \begin{cases} \dfrac{n - \dfrac{1}{2}}{n} \dfrac{\Gamma\left(\dfrac{(n-1)}{2}\right)}{\sqrt{\pi}\Gamma(\frac{n}{2})} \sum_{i=1}^{n-1} \sqrt{\dfrac{n-i}{i}} & \text{for } n \leq 340, \\[4ex] \dfrac{n - \dfrac{1}{2}}{n} \dfrac{1}{\sqrt{\dfrac{n\pi}{2}}} \sum_{i=1}^{n-1} \sqrt{\dfrac{n-i}{i}} & \text{for } n > 340, \end{cases}$$

where Γ is the Euler function. As pointed in Weron (2002), the Hurst exponent can be estimated more precisely as 0.5 plus the slope from the regression of $\frac{R}{S}(n) - E(\frac{R}{S}(n))$ regressed on $\log(n)$. This is referred to as *corrected empirical rescaled range analysis*.

3. Data

The data used in our analysis are daily Euro foreign exchange rates in 2018–2020. They are based on a regular daily concertation procedure between central banks across Europe and available by European Central Bank. In particular, we study the exchange rates for the following currencies: USD (United States dollar), JPY (Japanese yen), CZK (Czech koruna), DKK (Danish krone), GBP (Pound sterling), HUF (Hungarian forint), PLN (Polish złoty), RON (Romanian leu), SEK (Swedish krona), CHF (Swiss franc), ISK (Icelandic króna), NOK (Norwegian krone), HRK (Croatian kuna), RUB (Russian ruble), TRY (Turkish lira), AUD (Australian dollar), BRL (Brazilian real), CAD (Canadian dollar), CNY (Chinese yuan renminbi), HKD (Hong Kong dollar), IDR (Indonesian rupiah), ILS (Isreali shekel), INR (Indian rupee), KRW (South Korean won), MXN (Mexican peso), MYR (Malaysian ringgit), NZD (New Zealand dollar), PHP (Philippine peso), SGD (Singapore dollar), THB (Thai baht), and ZAR (South African rand).

In order to make the time series stationary, we follow a standard procedure of working with differences of logarithms of the rates. Figure 1 shows a selection of the data. We note that the volatility of the time series seems to be varying in time, which motivates us to use GARCH models for their modelling. Finding particular reasons for nonconstant volatility in the exchange rates data would need a standalone analysis. Here, we only note that this is not a new phenomenon; it has been studied in many papers (e.g., Feng et al., 2021; Kido, 2016; Manasseh et al., 2019; You & Liu, 2020; Zhou et al., 2020).
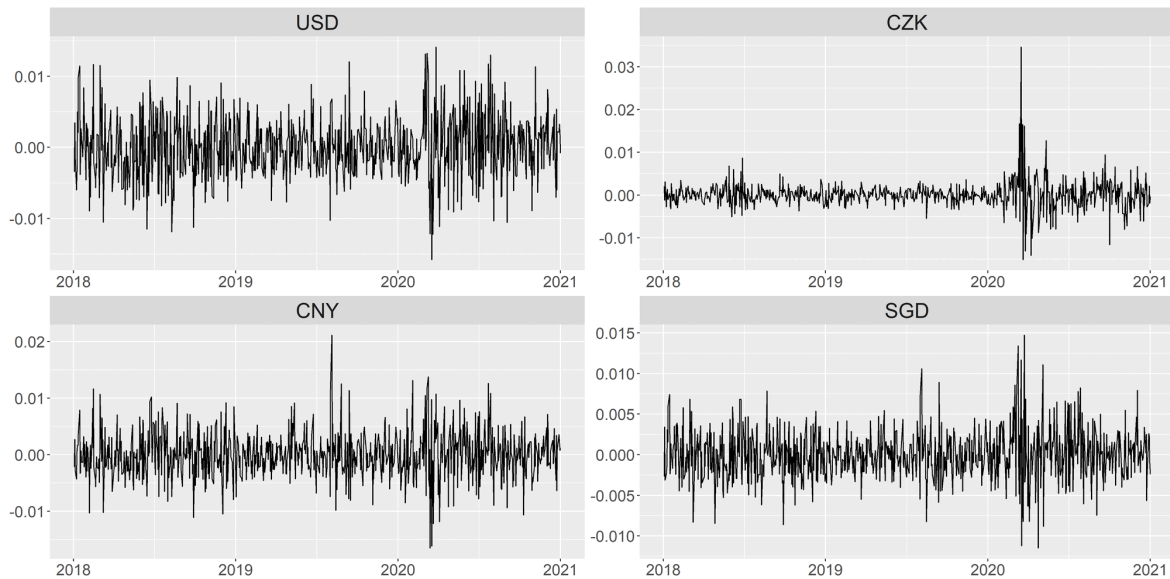
Figure 1: Sample of the data, differences of logarithms of the selected exchange rates

4. GARCH models

Let us recall that a standard autoregressive AR(p) model for stationary time serie x_t takes the form

$$x_t = \delta + a_1 x_{t-1} + \cdots + a_p x_{t-p} + u_t,$$

where the error term u is a white noise. The parameters are required to satisfy certain condition to ensure stationarity of the process (Kirchgässner et al., 2013). However, in financial applications it is often the case that the assumption of a constant variance of the

white noise is not consistent with observed data. The time varying variance of the data can be captured by GARCH processes, which model the variance σ_t^2 of the process u_t by the equation

$$\sigma_t^2 = \omega + \alpha_1 u_{t-1}^2 + \cdots + \alpha_p u_{t-p}^2 + \beta_1 \sigma_{t-1}^2 + \cdots + \beta_q \sigma_{t-q}^2,$$

where again the parameters are required to satisfy stationarity conditions. This process is known as GARCH(p,q) process; we refer the reader to Kirchgässner et al. (2013) for details.

We use `garchFit()` function from the R package `fGarch` (Wuertz et al., 2020) to estimate GARCH models and to obtain results of statistical tests necessary for evaluating the models. The residuals of the models are tested in order to assess the suitability of the proposed GARCH models. Following the standard procedures, implemented in the `fGarch` package, we use the Ljung-Box test for the residuals and the squared residuals and the heteroscedasticity test. In the model selection we consider autoregressive AR(p) processes with orders $p \leq 3$ with GARCH(p,q) error term with orders satisfying $p+q \leq 3$. From the models with residuals passing the tests given above on 5 % significance level, we select the model with the lowest Bayesian information criterion. Exchange rates IDR (Indonesian rupiah) and ILS (Isreali shekel) were excluded from the data due to fact that none of the model considered was suitable for them. The resulting models for the remaining exchange rates are given in Table 1.

Table 1: Autoregressive models with GARCH errors

| Model	Exchange rates
AR(0) + GARCH(1,1)	USD, DKK, HUF, JPY, GBP, SEK, RUB, AUD, CNY, CHF, BRL, HKD, INR, MXN, NZD, CAD, SGD, THB, ZAR
AR(0) + GARCH(2,1)	CZK
AR(0) + GARCH(1,2)	NOK, KRW, MYR, PHP
AR(1) + GARCH(1,1)	HRK, TRY
AR(2) + GARCH(1,1)	PLN
AR(3) + GARCH(1,2)	RON

5. Time varying Hurst exponents and their clustering

As outlined in the introduction, we use the approach from Cajueiro and Tabak (2004), Jiang et al. (2018), and Sensoy and Hacihasanoglu (2014), and we do not represent time series by a single estimated Hurst exponent. Instead, we represent it by sequence of Hurst exponents estimated from shorter time windows to capture regime changes within data. The main reason is that in many financial time series we can observe cycles of irregular length in which the dynamics varies. It is reasonable to assume that this would be also true even for Hurst exponent. Another reason for choosing a sequence of Hurst exponents over one particular Hurst exponent estimate would be that we might be also interested in studying reaction of exchange rates dynamics during specific time window on information that were dominating through the specific time.

We choose a rolling window approach to estimate sequence of Hurst exponent for each exchange rate with window size selected to be 252 days, which is approximately one year of data (since the data are available only on business days). This means that for each sequence

$\{X_j\}_{j=i}^{i+w-1}$ with w being size of window and $i = 1, \ldots, n - w + 1$, we estimated H_i as Hurst exponent for particular time period. This results in sequence of Hurst exponents $\{H_i\}_{i=1}^{n-w+1}$.

For determining clusters, hierarchical clustering was employed with Ward's minimum variance method using function `hclust()` from the `stats` R package. The distance was chosen as squared Euclidean distance between vectors of time-varying Hurst exponents which is required due to usage of the Ward's algorithm. We note that a popular similarity measure based on correlations is not applicable here. Two evolutions of Hurst exponents, which differ by a constant, have a perfect correlation. However, they might be on opposite sides of $H = 1/2$ and thus exhibiting different characteristics, which we would like to take into account. To determine the number of clusters, we used silhouette criterion (Rousseeuw, 1987) using function `silhouette()` from `cluster` R package (Maechler et al., 2019).

We employed four different calculations of Hurst exponent, as described in Section 2, resulting in four different vectors of time-varying Hurst exponent for each exchange rate. Examples of the time-varying Hurst exponents are shown in Figure 2. As can be seen, time-varying Hurst exponents for particular exchange rate significantly differs by used estimation technique so it is meaningful to carry out cluster analysis for every one of them. Thus, resulting in 4 clusterings of exchange rate market. Dendrograms and the resulting clusters are presented in Figures 3–5.
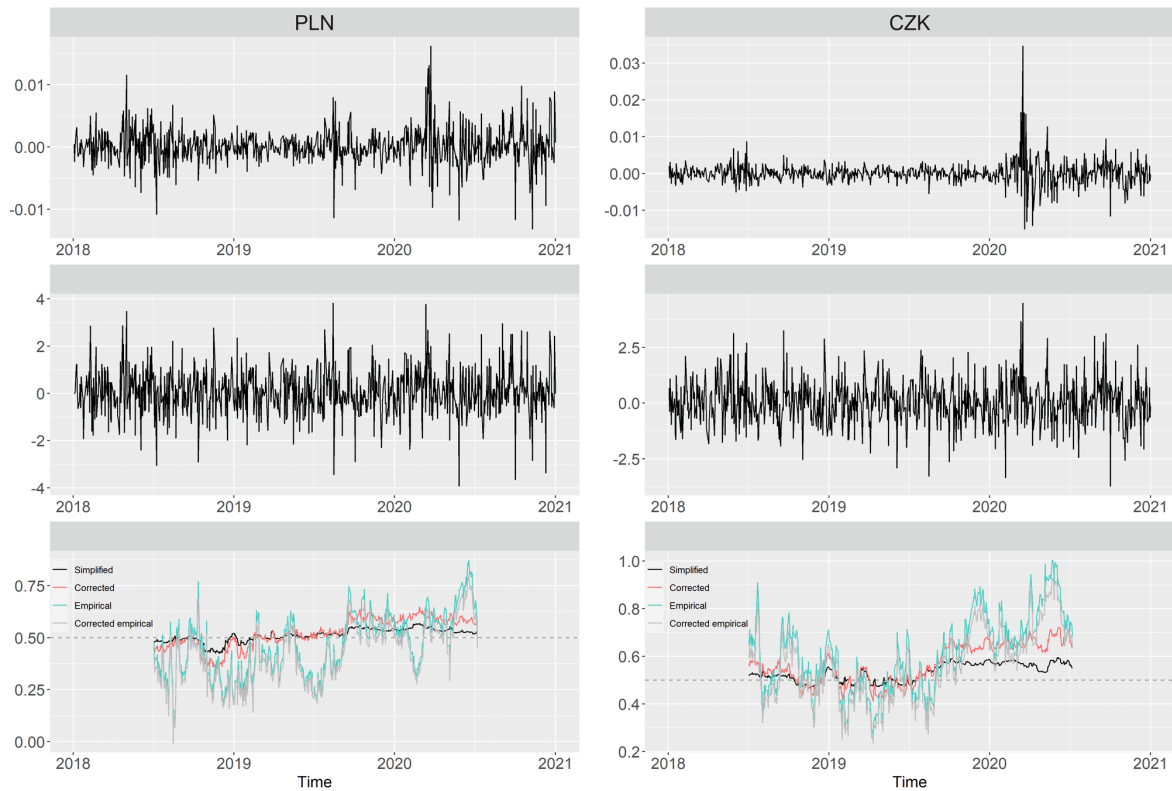

Figure 2: Time dependent estimates of the Hurst exponents for selected currencies (bottom figures), together with the original data (top) and standardized residuals (middle)

6. Comparison of clusterings and final clusters

Clusterings presented in the previous section are not identical; however, in the case of certain pairs of exchange rates, they were in the same cluster in all four clusterings. We
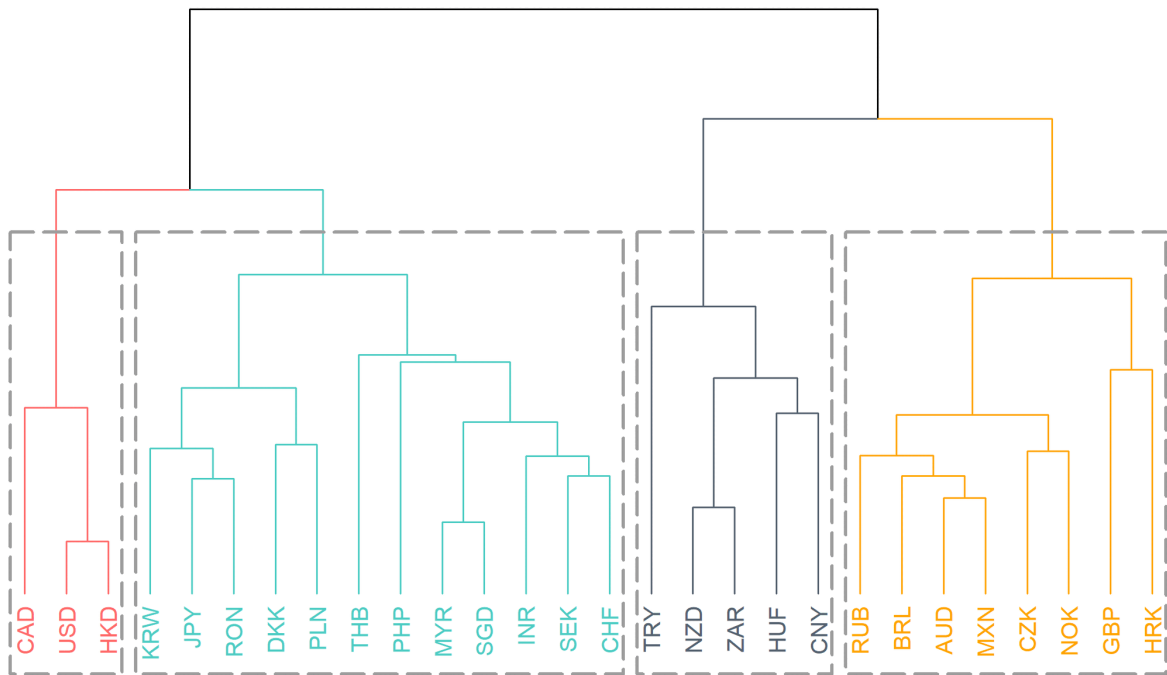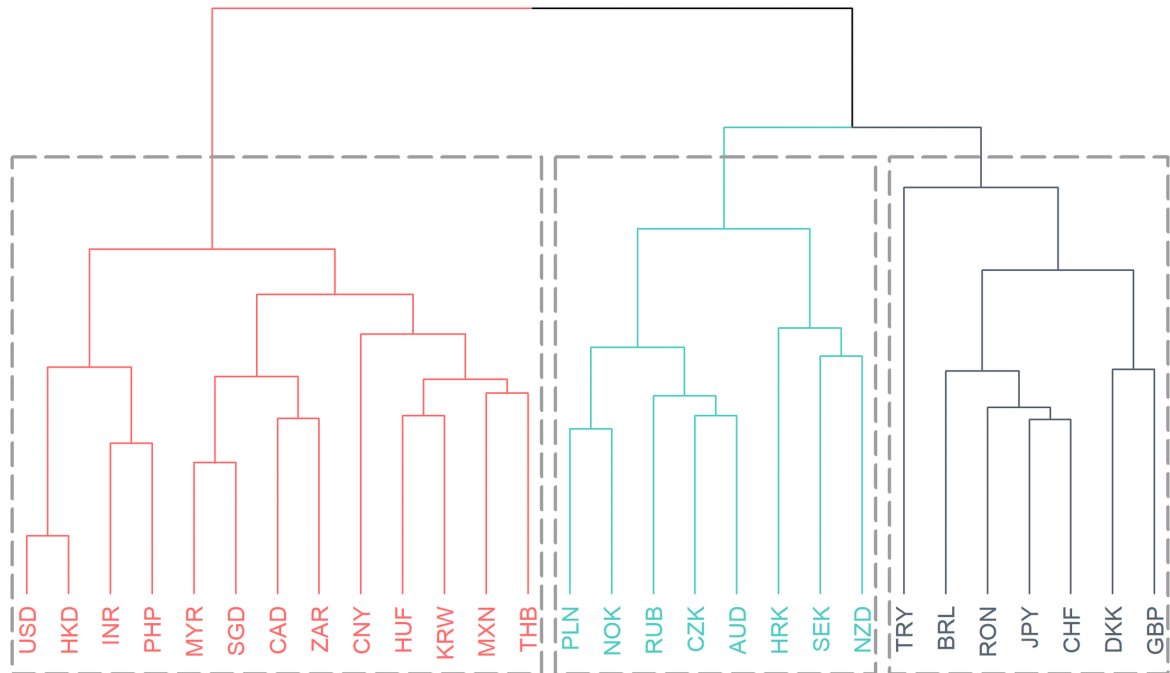
Figure 3: Hierarchical clustering of the currencies based on simplified Hurst exponent clustered by Ward's algorithm using squared Euclidean distance as similarity measure. Optimal clusters selected via silhouette criterion are visualised by the dashed frames.
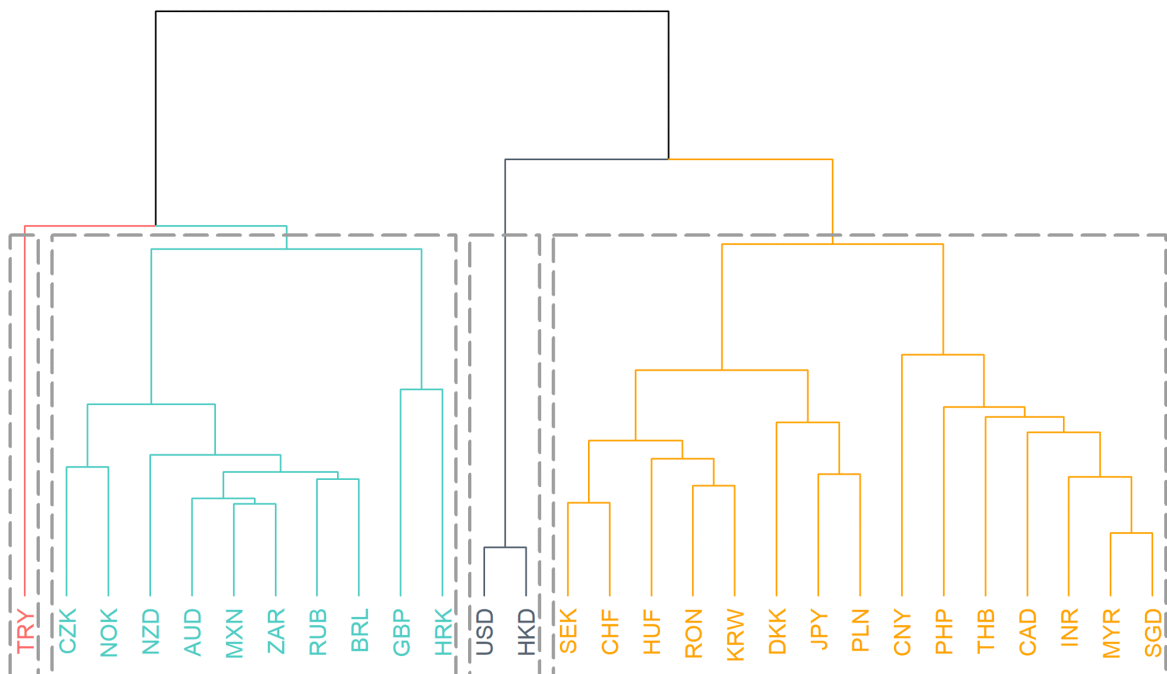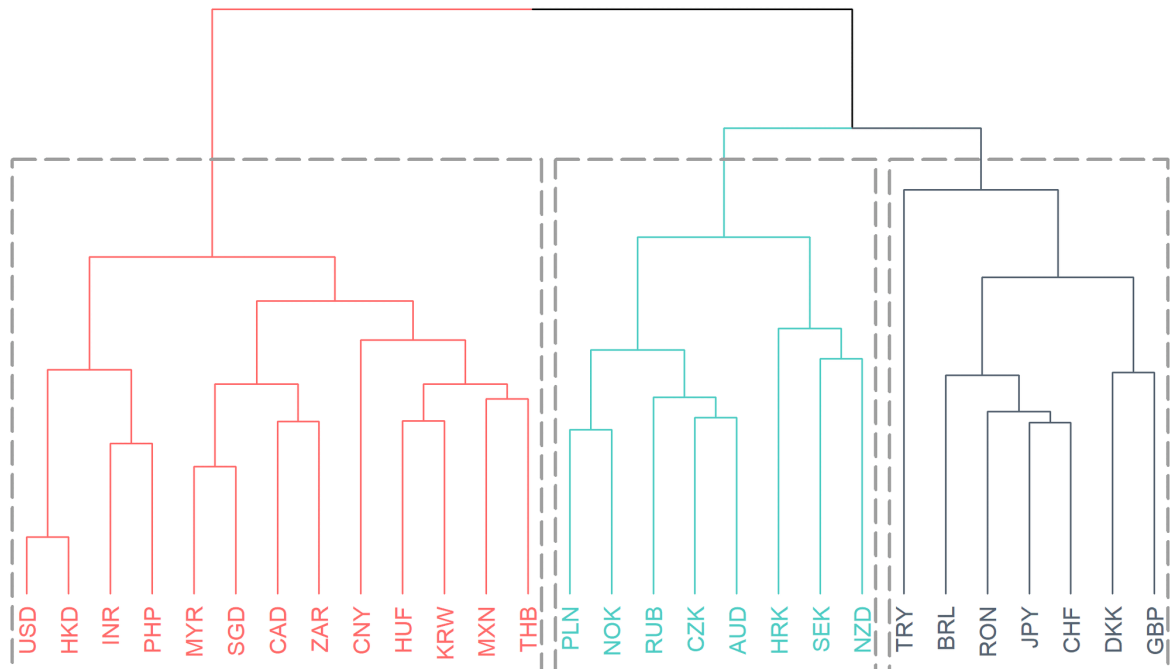

Figure 4: Hierarchical clustering of the currencies based on corrected Hurst exponent clustered by Ward's algorithm using squared Euclidean distance as similarity measure. Optimal clusters selected via silhouette criterion are visualised by the dashed frames.

Figure 5: Hierarchical clustering of the currencies based on empirical Hurst exponent clustered by Ward's algorithm using squared Euclidean distance as similarity measure. Optimal clusters selected via silhouette criterion are visualised by the dashed frames.

Figure 6: Hierarchical clustering of the currencies based on corrected empirical Hurst exponent clustered by Ward's algorithm using squared Euclidean distance as similarity measure. Optimal clusters selected via silhouette criterion are visualised by the dashed frames.

consider this to be a strong indicator that the Hurst exponent has a similar evolution for these two rates. As the number of such cases decreases, also the similarity can be seen as weaker. Naturally, in many cases, the given pair of the rates was never in the same cluster.

We associate the clustering results with a network, whose nodes are the exchange rates. Two nodes are connected by an edge, if they were in the same cluster at least once. The weight of the edge is given by the number of such clustering. The resulting network is shown in Figure 7.

If we consider the edges with the weight above a certain threshold, the network splits into several connected components. The nodes in these components are therefore representing sets of exchange rates, for which the evolution of the Hurst exponent is similar. Therefore, we take the connected components as the final clusters of our analysis. The choice of the threshold is subjective, we base it on visualizing the networks corresponding to different thresholds. Depending of the data, we might need to find a trade-off between a large number of small components with strong connections between the nodes, and a small number of large components with weaker ties. In our particular case we compare the components emerging from the thresholds 4 (the maximum possible weight of an edge) and 3 (which means that the exchange rates have to be in the same clusters at least 3 times out of 4, in order to be connected by an edge in the network). We do not consider lower values for a threshold; requiring the edge of the weight to be at least 3 means that the nodes connected by an edge must be in the same cluster in more than half of the cases. The clusters consisting of more than one node are presented in Tables 2 and 3.

Both clusterings seem reasonable. We can identify nodes in the clusters which can be expected to be in the same cluster based on the dependence of the economies and financial markets in the given countries. In the network from the threshold 4, we see a small cluster containing United States dollar and Hong Kong dollar. A larger cluster, containing six nodes, includes exchange rates of currencies in countries located in the south, southeast and east Asia - Indian rupee, South Korean won, Malaysian ringgit, Philippine peso, Singapore dollar, Thai baht. We note, however, that being in the same cluster does not mean a similar evolution of the exchange rate itself. Instead, it means a similar evolution of the Hurst exponent. Therefore, a more detailed interpretation of the clusters would need a more careful on the factors that might influence this feature of the exchange rates.

Table 2: Final clustering of the exchange rates for the threshold 4—nontrivial clusters (containing more than one exchange rate)

| Cluster	Exchange rates
1	USD, HKD
2	JPY, DKK, RON, CHF
3	CZK, NOK, HRK, RUB, AUD
4	GBP, BRL
5	HUF, CNY
6	PLN, SEK
7	INR, KRW, MYR, PHP, SGD, TBH
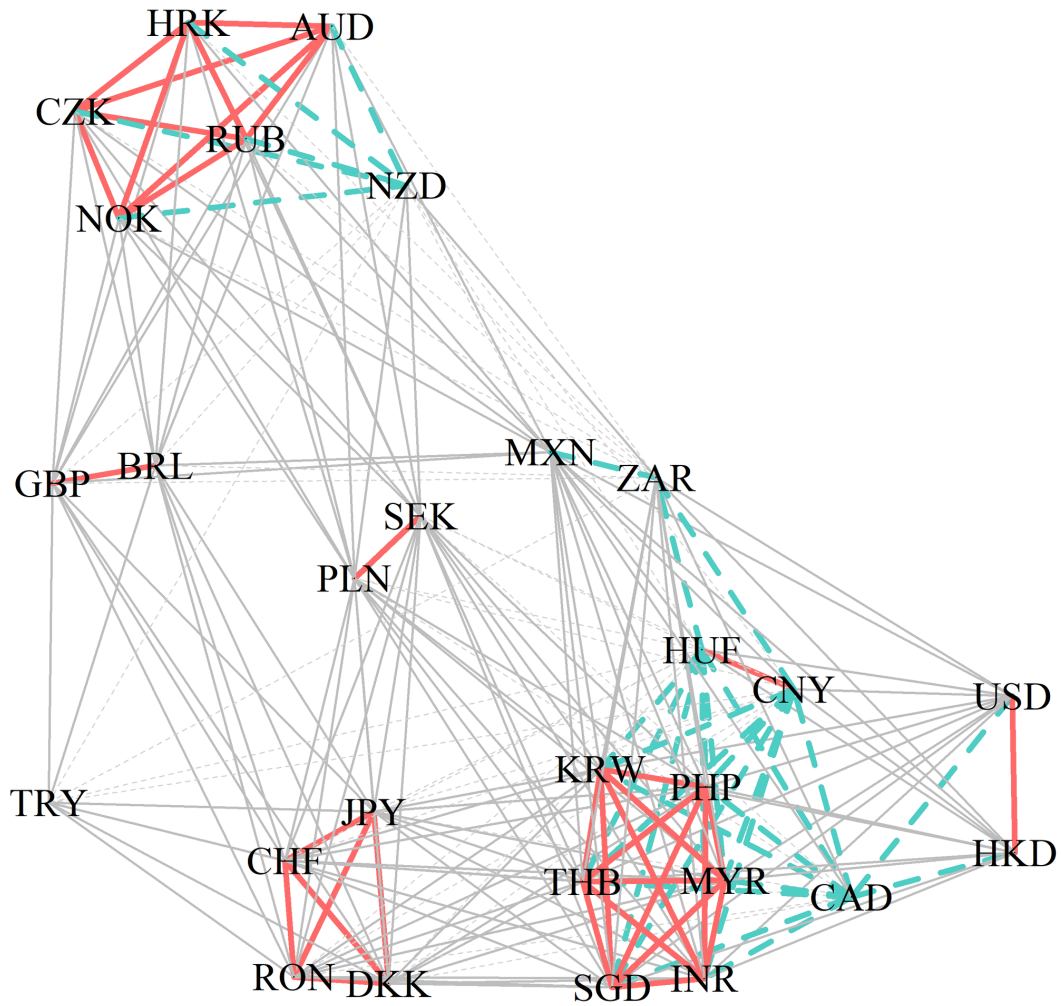
Figure 7: Network constructed from the clusterings in Figures 3–6 with the vertices corresponding to exchange rates and the weight of the edges corresponding to number of times each pair of exchange rates ended up together in a cluster. The weight of the edge is visualized by the width of the line, the type of the line and by its colour (1 = thin, dashed, light-grey, 2 = thin, solid, grey, 3 = thick, dashed, green, 4 = thick, solid, red).

Table 3: Final clustering of the exchange rates for the threshold 3—nontrivial clusters (containing more than one exchange rate)

Cluster	Exchange rates
1	USD, HUF, CAD, CNY, HKD, INR, KRW, MXN, MYR, PHP, SGD, THB, ZAR
2	JPY, DKK, RON, CHF
3	CZK, NOK, HRK, RUB, AUD, NZD
4	GBP, BRL
5	PLN, SEK

7. Conclusions

Many financial time series exhibit long-range dependence. We used this property for clustering the time series based on Hurst exponent which measures this dependence. We
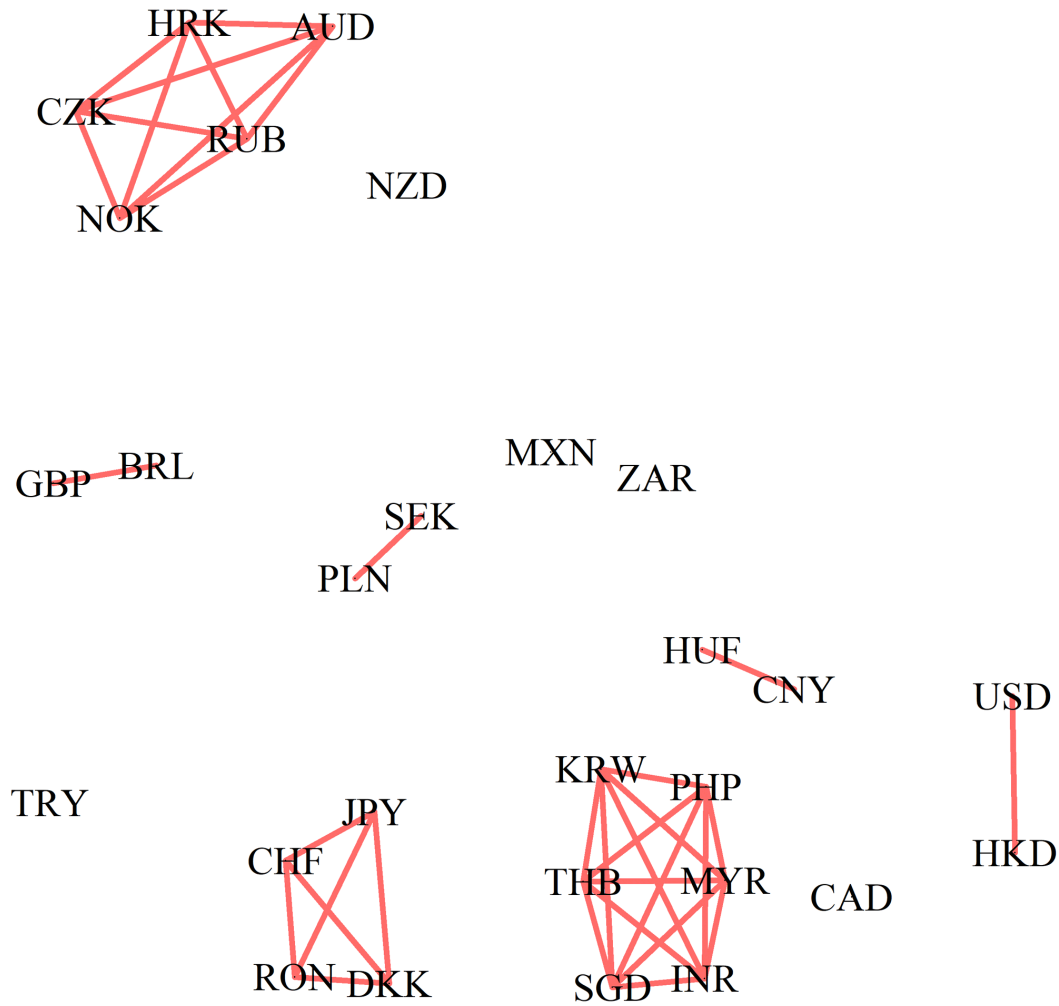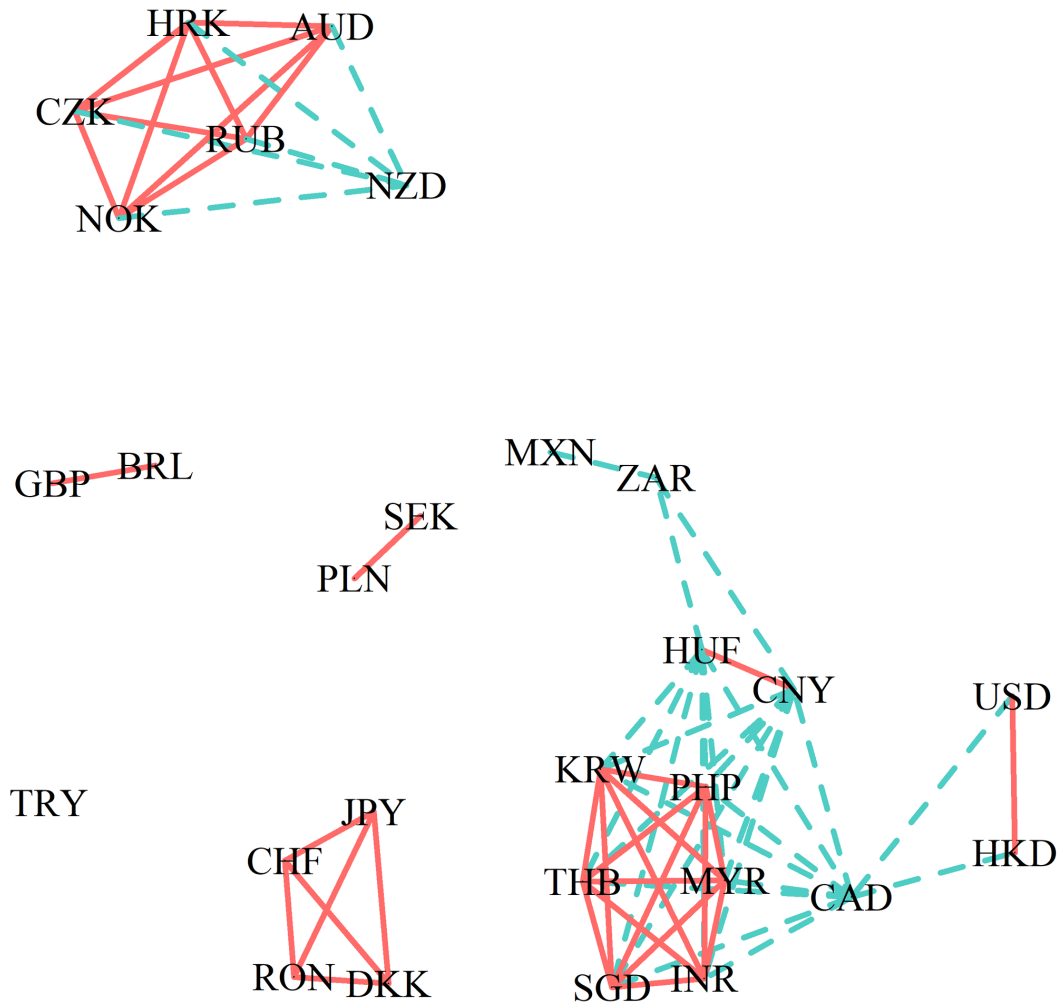
Figure 8: Network obtained from Figure 7 by only considering edges with weights equal to 4. The vertices correspond to exchange rates and the weights of the edges correspond to number of times each pair of exchange rates ended up together in a cluster.

proposed a clustering procedure which uses several different estimates of the Hurst exponent and clusters their values obtained by a rolling window method. As a final step of our procedure, clusterings originating from individual methods of the Hurst exponent were compared. In our example of exchange rates, it turns out that we are able to create final clustering by requiring that members of each cluster are in the same cluster at least the specified number of times in the individual clusterings. We expect the same to hold also in the case of other data since "similar time series" should appear in the same cluster often, when considering different details of clustering procedure. Therefore, our approach can be directly applied also to other time series.

The results which we have obtained provide a new application of rolling window approach to Hurst exponent estimation, used earlier in Cajueiro and Tabak (2004), Jiang et al. (2018), and Sensoy and Hacihasanoglu (2014). Moreover, they make it possible to extend other clustering analyses such as Lahmiri (2016), by allowing to use more than one time criterion (e.g., several estimation methods in our particular case).

The extension of our results can go in two directions. The first one consists of a more detailed interpretation of the clustering. As we noted, the estimates of the Hurst exponent provide an information about the underlying time series and we might study the external

Figure 9: Network obtained from 7 by only considering edges with weights greater than 2. The vertices correspond to exchange rates and the weights of the edges correspond to number of times each pair of exchange rates ended up together in a cluster. The weight of the edge is visualized by the type of the line and by its colour (3 = dashed, green, 4 = solid, red).

factors which lead to this behaviour of the data. This might also give a better understanding of clusters and why certain exchange rates (or other data considered) appear in the same or in different clusters, respectively.

The other direction involves using different methods to construct individual clusterings. The final comparison of clusters does not have any limitation on the number of clusterings which enter it, neither on methods used to obtain them. There are many methods for estimating Hurst exponents, other distances between vectors of Hurst exponents may be considered, we may use different clustering methods. Individual clusterings might get different weights and instead of counting the number of occurrences in the same cluster, it is possible to weight them. It might be insightful to see how these different approaches influence the final clusters.

A possible limitation might be the need of finding a suitable trade-off between clearly distinguished components in the network and the number of isolated nodes, corresponding to clusters containing one time series. If the condition for the existence of an edge between nodes is not sufficiently strict, i.e., only a small number of occurrences in the same cluster is required, the components are often large, which may not be always desirable. On the other hand, a high threshold often leaves a lot of nodes without an edge, leading to one-

element clusters. However, we may be interested in finding similar time series to most of the data, instead of concluding that they form a separate cluster. A possible solution might be a modification of our final clustering step. Instead of considering the components of the network, various methods for finding so called communities in connected networks can be employed. They aim to divide the nodes into communities, which are characterized by many edges within the nodes in a community and a small number of edges between nodes in different communities. Reviews of such methods can be found in Fortunato (2010) and Javed et al. (2018). The proposed method and its possible modifications outlined above provide a new approach for clustering time series using networks and communities, considered in Ferreira and Zhao (2015).

To conclude, we note again that the proposed approach can be used to analyze any time series with long-range dependence, or time series for which their regimes—persistent, anti-persistent or having quickly decaying correlations—need to be distinguished. Therefore we consider it to be an interesting addition to the topic of clustering time series with these properties.

Acknowledgment

References

Aghabozorgi, S., Shirkhorshidi, A. S., & Wah, T. Y. (2015). Time-series clustering–a decade review. *Information Systems*, *53*, 16–38. https://doi.org/10.1016/j.is.2015.04.007

Annis, A. A., & Lloyd, E. H. (1976). The expected value of the adjusted rescaled hurst range of independent normal summands. *Biometrika*, *63*(1), 111–116. https://doi.org/10.1093/biomet/63.1.111

Auer, B. R. (2016). Pure return persistence, hurst exponents and hedge fund selection – a practical note. *Journal of Asset Management*, *17*(5), 319–330. https://doi.org/10.1057/jam.2016.7

Beran, J. (2017). *Statistics for long-memory processes*. Routledge. https://doi.org/10.1201/9780203738481

Borchers, H. W. (2019). *pracma: Practical numerical math functions* (Version 2.2.9) [Computer software]. The Comprehensive R Archive Network. https://cran.r-project.org/package=pracma

Cajueiro, D. O., & Tabak, B. M. (2004). The hurst exponent over time: Testing the assertion that emerging markets are becoming more efficient. *Physica A: Statistical Mechanics and its Applications*, *336*(3-4), 521–537. https://doi.org/10.1016/j.physa.2003.12.031

Cajueiro, D. O., & Tabak, B. M. (2009). Testing for long-range dependence in the Brazilian term structure of interest rates. *Chaos, Solitons & Fractals*, *40*(4), 1559–1573. https://doi.org/10.1016/j.chaos.2007.09.054

Di Matteo, T. (2007). Multi-scaling in finance. *Quantitative Finance*, *7*(1), 21–36. https://doi.org/10.1080/14697680600969727

Feng, G.-F., Yang, H.-C., Gong, Q., & Chang, C.-P. (2021). What is the exchange rate volatility response to COVID-19 and government interventions? *Economic Analysis and Policy*, *69*, 705–719. https://doi.org/10.1016/j.eap.2021.01.018

Ferreira, L. N., & Zhao, L. (2015). A time series clustering technique based on community detection in networks. *Procedia Computer Science*, *53*, 183–190. https://doi.org/10.1016/j.procs.2015.07.293

Fortunato, S. (2010). Community detection in graphs. *Physics Reports*, *486*(3-5), 75–174. https://doi.org/10.1016/j.physrep.2009.11.002

Fu, T.-c. (2011). A review on time series data mining. *Engineering Applications of Artificial Intelligence*, *24*(1), 164–181. https://doi.org/10.1016/j.engappai.2010.09.007

Han, J., & Ge, Z. (2020). Effect of dimensionality reduction on stock selection with cluster analysis in different market situations. *Expert Systems with Applications*, *147*, 113226. https://doi.org/10.1016/j.eswa.2020.113226

Hurst, H. E. (1951). Long-term storage capacity of reservoirs. *Transactions of the American Society of Civil Engineers*, *116*(1), 770–799. https://doi.org/10.1061/taceat.0006518

Iorio, C., Frasso, G., D'Ambrosio, A., & Siciliano, R. (2018). A p-spline based clustering approach for portfolio selection. *Expert Systems with Applications*, *95*, 88–103. https://doi.org/10.1016/j.eswa.2017.11.031

Javed, M. A., Younis, M. S., Latif, S., Qadir, J., & Baig, A. (2018). Community detection in networks: A multidisciplinary review. *Journal of Network and Computer Applications*, *108*, 87–111. https://doi.org/10.1016/j.jnca.2018.02.011

Jiang, Y., Nie, H., & Ruan, W. (2018). Time-varying long-term memory in bitcoin market. *Finance Research Letters*, *25*, 280–284. https://doi.org/10.1016/j.frl.2017.12.009

Kido, Y. (2016). On the link between the US economic policy uncertainty and exchange rates. *Economics Letters*, *144*, 49–52. https://doi.org/10.1016/j.econlet.2016.04.022

Kirchgässner, G., Wolters, J., & Hassler, U. (2013). *Introduction to modern time series analysis*. Springer. https://doi.org/10.1007/978-3-642-33436-8

Lahmiri, S. (2016). Clustering of Casablanca stock market based on hurst exponent estimates. *Physica A: Statistical Mechanics and its Applications*, *456*, 310–318. https://doi.org/10.1016/j.physa.2016.03.069

Maechler, M., Rousseeuw, P., Struyf, A., Hubert, M., & Hornik, K. (2019). *pracma: Practical numerical math functions* (Version 2.1.0) [Computer software]. The Comprehensive R Archive Network. https://cran.r-project.org/package=cluster

Maharaj, E. A., D'Urso, P., & Caiado, J. (2019). *Time series clustering and classification*. Chapman; Hall/CRC. https://doi.org/10.1201/9780429058264

Manasseh, C. O., Chukwu, N. O., Abada, F. C., Ogbuabor, J. E., Lawal, A. I., & Alio, F. C. (2019). Interactions between stock prices and exchange rates: An application of multivariate VAR-GARCH model. *Cogent Economics & Finance*, *7*(1), 1681573. https://doi.org/10.1080/23322039.2019.1681573

Mandelbrot, B. B. (1975). Limit theorems on the self-normalized range for weakly and strongly dependent processes. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, *31*(4), 271–285. https://doi.org/10.1007/bf00532867

Mandelbrot, B. B., & Wallis, J. R. (1969). Robustness of the rescaled range R/S in the measurement of noncyclic long run statistical dependence. *Water Resources Research*, *5*(5), 967–988. https://doi.org/10.1029/wr005i005p00967

Massahi, M., Mahootchi, M., & Arshadi Khamseh, A. (2020). Development of an efficient cluster-based portfolio optimization model under realistic market conditions. *Empirical Economics*, *59*(5), 2423–2442. https://doi.org/10.1007/s00181-019-01802-5

O'Connell, P., Koutsoyiannis, D., Lins, H. F., Markonis, Y., Montanari, A., & Cohn, T. (2016). The scientific legacy of Harold Edwin hurst (1880–1978). *Hydrological Sciences Journal*, *61*(9), 1571–1590. https://doi.org/10.1080/02626667.2015.1125998

Peters, E. E. (1994). *Fractal market analysis: Applying chaos theory to investment and economics*. John Wiley & Sons.

Rousseeuw, P. J. (1987). Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, *20*, 53–65. https://doi.org/10.1016/0377-0427(87)90125-7

Sensoy, A., & Hacihasanoglu, E. (2014). Time-varying long range dependence in energy futures markets. *Energy Economics*, *46*, 318–327. https://doi.org/10.1016/j.eneco.2014.09.023

Song, J. Y., Chang, W., & Song, J. W. (2019). Cluster analysis on the structure of the cryptocurrency market via Bitcoin–Ethereum filtering. *Physica A: Statistical Mechanics and its Applications*, *527*, 121339. https://doi.org/10.1016/j.physa.2019.121339

Taqqu, M. S., Teverovsky, V., & Willinger, W. (1995). Estimators for long-range dependence: An empirical study. *Fractals. Complex Geometry, Patterns, and Scaling in Nature and Society*, *03*(04), 785–798. https://doi.org/10.1142/s0218348x95000692

Weron, R. (2002). Estimating long-range dependence: Finite sample properties and confidence intervals. *Physica A: Statistical Mechanics and its Applications*, *312*(1-2), 285–299. https://doi.org/10.1016/s0378-4371(02)00961-5

Wuertz, D., Setz, T., Chalabi, Y., Boudt, C., Chausse, P., & Miklovac, M. (2020). *fGarch: Rmetrics - autoregressive conditional heteroskedastic modelling* (Version 3042.83.2) [Computer software]. The Comprehensive R Archive Network. https://cran.r-project.org/package=fGarch

You, Y., & Liu, X. (2020). Forecasting short-run exchange rate volatility with monetary fundamentals: A GARCH-MIDAS approach. *Journal of Banking & Finance*, *116*, 105849. https://doi.org/10.1016/j.jbankfin.2020.105849

Zhou, Z., Fu, Z., Jiang, Y., Zeng, X., & Lin, L. (2020). Can economic policy uncertainty predict exchange rate volatility? New evidence from the GARCH-MIDAS model. *Finance Research Letters*, *34*, 101258. https://doi.org/10.1016/j.frl.2019.08.006