Accepted Manuscript

Dynamic correlations at different time-scales with empirical mode decomposition

Noemi Nava, T. Di Matteo, Tomaso Aste

PII: S0378-4371(18)30198-5
DOI: https://doi.org/10.1016/j.physa.2018.02.108
Reference: PHYSA 19228

To appear in: *Physica A*

Received date : 22 August 2017
Revised date : 22 December 2017

Please cite this article as: N. Nava, T. Di Matteo, T. Aste, Dynamic correlations at different time-scales with empirical mode decomposition, *Physica A* (2018), https://doi.org/10.1016/j.physa.2018.02.108

***Highlights (for review)**

Highlights

- Investigation of correlations at different time horizons by combining Empirical Mode Decomposition and Pearson Cross-correlations
- Novel and simple methodology with effective outcomes for correlations at different time horizons
- Demonstration for S&P500, VIX and IPC indices

Dynamic correlations at different time-scales with Empirical Mode Decomposition

Noemi Nava[a,b], T. Di Matteo[a,b,c,d], Tomaso Aste[a,b]

[a] *Department of Computer Science, University College London, Gower Street, London, WC1E 6BT, UK*
[b] *Systemic Risk Centre, London School of Economics and Political Sciences, London, WC2A2AE, UK*
[c] *Department of Mathematics, King's College London, The Strand, London, WC2R 2LS, UK*
[d] *Complexity Science Hub Josefstaedter Strasse 39, A 1080 Vienna, Austria.*

Abstract

We introduce a simple approach which combines Empirical Mode Decomposition (EMD) and Pearson's cross-correlations over rolling windows to quantify dynamic dependency at different time scales. The EMD is a tool to separate time series into implicit components which oscillate at different time-scales. We apply this decomposition to intraday time series of the following three financial indices: the S&P 500 (USA), the IPC (Mexico) and the VIX (volatility index USA), obtaining time-varying multidimensional cross-correlations at different time-scales. The correlations computed over a rolling window are compared across the three indices, across the components at different time-scales and across different time lags. We uncover a rich heterogeneity of interactions, which depends on the time-scale and has important led-lag relations that could have practical use for portfolio management, risk estimation and investment decisions.

Keywords: Time-scale-dependent correlation, time-dependent correlation, empirical mode decomposition

Email address: t.aste@ucl.ac.uk (Tomaso Aste)

1. Correlation structures in financial time series

Financial time series are correlated, and the structure of these correlations reflects market properties [1, 2, 3, 4]. Analysis of correlations is important for the construction of diversified portfolios, to reduce investment risks and to orientate investment decisions. The correlation structure of financial markets has been one of the major domain of studies in the last couples of decades [5]. Researches have highlighted that the correlation structure is associated with industry sectors, it is time dependent and it provides a dynamical map of market states which it follows seasonalities, trends and has abrupt changes during market instabilities [6, 7, 2, 8, 9, 10, 11, 12, 13, 14].

Financial markets operate at different time horizons [15], and characterizing the relation between market prices at different time-scales is essential to capture the complexity of market dynamics for portfolio management, risk management and investments [16, 17, 18, 19, 20, 21, 22]. It is well established and documented that correlations between stock returns vary over time (see for instance [2, 23, 24, 25]). It is instead less understood and established how correlations between financial assets vary over time-scales [26, 6]. Most studies focus on the analysis of a specific time-scale. However, changes of correlation at different time-scales have important practical consequences. For instance, if the correlation between two assets varies across time-scales, then market participants with short and long term-horizons have different risk exposures and must adapt their strategies according to the different parts of the correlation spectrum. Furthermore, investigating both the time dependent and the time-scale dependent dynamics of correlations can provide insights on the collective behaviour of traders with different strategies [23, 26]. This is the topic of the present paper where we use a simple methodology to perform this research.

In the literature several methodologies have been proposed to analyse correlations and dependencies between time series at different frequencies. Wavelets and the wavelet transform modulus maxima [27] are among the first and most used methodologies to simultaneously analyse the correlation and multiscale structures [28, 25, 24]. Another noticeable method is the multivariate extension of the multifractal detrended fluctuation analysis, which combines established methods for multifractal characterisation of signals with quantification of correlations by using local piecewise polynomial detrending

[29, 30, 31].

The approach in this paper is a simple combination of the Pearson cross-correlation with the decomposition of a time series at different time-scales obtained with the Empirical Mode Decomposition [32]. The EMD method decomposes the signal into a finite set of nearly orthogonal oscillating components, called Intrinsic Mode Functions (IMF), and a residue. The obtained components are retrieved from the local maxima and the local minima of the data without imposing any kind of functional form or oscillation frequency. This is an important difference with respect to the wavelet transform (see [24, 25]) which instead requires to use a specific a priori filter function [33]. Therefore, the EMD relies on less assumptions, it is a fully data-driven decomposition which can be applied to non-stationary and non-linear data [32].

Our approach is similar to the methodology recently introduced by Chen et al. [34], who proposed to use the Empirical Mode Decomposition to estimate the so-called Time-Dependent Intrinsic Correlation (TDIC). In this approach, two time series are first decomposed into a set of IMFs oscillating at different time scales. Then, the Pearson correlation is calculated in an adaptive window whose length depends on the instantaneous period of the IMFs. In this paper, we introduce a simplified version of this approach. Simplification consists of two main elements: first, we adopt a single size window for each IMF; second, when decomposing the time series, we do not add white noise. We verified that these simplifications do not affect significantly our results providing however a more parsimonious approach that requires a smaller numbers of parameters and assumptions.

We apply our methodology to intra-day data (30 seconds) of three indices: the S&P 500 (USA), the IPC (Mexico) and the VIX (volatility index USA). We compute cross-correlations and lagged cross-correlations between the different IMFs. This yields to dynamic cross-correlations across time scales. The results uncover the presence of cross-scale coupling between the time series and identify some relevant led-lag relations at specific time-scales which could be relevant for practical purposes in portfolio management.

This paper is organized as follows. In Section 2, we introduce the basic concepts of the EMD and the IMF. In Section 3, the computation of cross-correlations across time-scales, time-lags and time-windows is described. Section 4 reports the application to real data on three indices: the S&P 500 (USA), the IPC (Mexico) and the VIX. The discussions and conclusions are provided in Section 5.

2. Empirical mode decomposition (EMD)

The EMD method identifies a finite set of oscillations with scale defined by the local maxima and the local minima of the data itself. Each oscillation is empirically derived from the data and is referred to as an Intrinsic Mode Function (IMF). An IMF must satisfy two criteria [32]:

1. The number of extrema and the number of zero crossings must either be equal or differ at most by one;
2. At any point, the mean value of the envelope defined by the local maxima and the envelope defined by the local minima is zero.

The first condition forces an IMF to be a narrow-band signal with no riding waves. The second condition ensures that the instantaneous time-scale will not have fluctuations arising from an asymmetric wave form [32]. The IMFs are obtained through a process called sifting process which uses local extrema to separate oscillations starting with the highest time-scale. Given a time series $X(t)$, $t = 1, 2, ..., T$, the process decomposes it into a finite number of components, the $\text{IMF}_k(t)$, $k = 1, ..., n$, and a residue $r_n(t)$. The residue is the non-oscillating drift of the data. If the decomposed data consists of uniform scales in the time-scale space, the EMD acts as a dyadic filter and the total number of IMFs is approximately equal to n= $\log_2(T)$ [35]. At the end of the decomposition process, the original time series can be reconstructed as:

$$X(t) = \sum_{k=1}^{n} \text{IMF}_k(t) + r_n(t). \tag{1}$$

The EMD is implemented through the following steps [32]:

1. Initialize the residue to the original time series $r_0(t) = X(t)$ and set the IMF index $k = 1$.
2. Extract the k^{th} IMF:
   (a) initialize $h_0(t) = r_{k-1}(t)$ and set the iteration counter $i = 1$;
   (b) find the local maxima and the local minima of $h_{i-1}(t)$ (see Figure 1(a));
   (c) create the upper envelope $E_u(t)$ by interpolating between the local maxima and, analogously, create lower envelope $E_l(t)$ by interpolating the local minima (see Figure 1(b));

(d) calculate the mean of both envelopes as $m_{i-1}(t) = \frac{E_u(t)+E_l(t)}{2}$ (see Figure 1(c));
(e) subtract the mean envelope from the input time series, obtaining $h_i(t) = h_{i-1}(t) - m_{i-1}(t)$, see Figure 1(d);
(f) verify if $h_i(t)$ satisfies the IMF's conditions:
- if $h_i(t)$ does not satisfy the IMF's conditions, increase $i = i+1$ and repeat the sifting process from step (b) (see Figure 1(d));
- if $h_i(t)$ satisfies the IMF's conditions, set $\text{IMF}_k(t) = h_i$ and define $r_k(t) = r_{k-1}(t) - \text{IMF}_k(t)$ (see Figure 2).

3. When the residue $r_k(t)$ is either a constant, a monotonic slope or contains only one extrema stop the process, otherwise continue the decomposition from step 2, setting $k = k + 1$.

In Figure 1, we exemplify some steps of the sifting process. After one iteration of the sifting process, the function $h_1(t)$ is obtained (Figure 1(d)). In this example, the resulting function is not symmetric and does not have zero mean, hence it is not an IMF yet. Consequently, more iterations of the sifting processes need to be applied to extract the first IMF of the input time series. These further iterations are shown in Figures 2(a), 2(b) and 2(c) with the last sifting iteration which extracts the first IMF, shown in Figure 2(d).

The sifting process eliminates the riding waves and smooths uneven amplitudes [32]. This process terminates when the local mean of the extracted IMF is zero. The difficulty is that this condition can only be approximated and in order to avoid over-sifting and converting meaningful IMFs into meaningless fluctuations with constant amplitude, a stopping criterion needs to be implemented.

It must be noted that the EMD is based on the timescale separation and does not impose orthogonality, implying that in general the sum of the variance of the components and the residue differs from the variance of the input time series. Although the non-orthogonality is conceptually of great relevance, in most practical cases the overlap between components is small and the difference between the sum of the variance of the components and the total variance remains small [32].

3. Cross-correlations on IMF

Let us consider two time series $X(t)$ and $Y(t)$ with $t = 1, 2, \ldots, T$, with equal length T and with equal intervals of time s between observations.

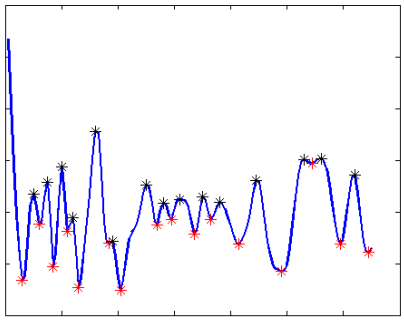

(a) Local maxima and minima.

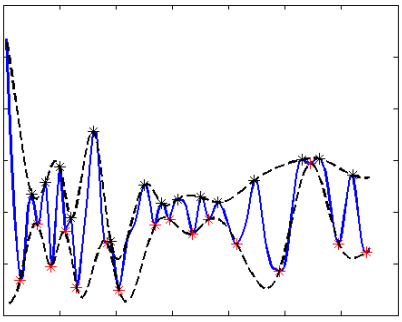

(b) Upper and lower envelopes.

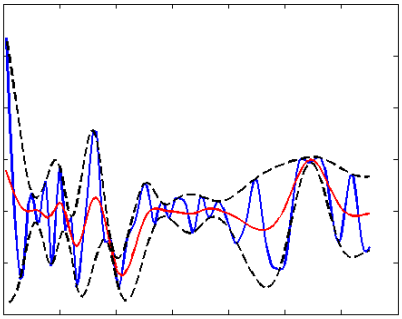

(c) Envelope mean.

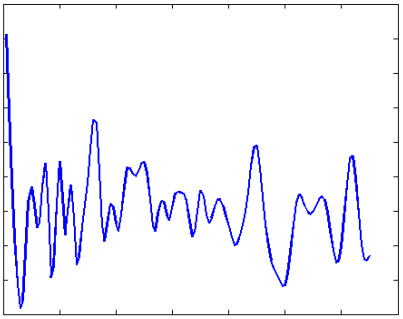

(d) Time series after one sifting step.

Figure 1: Example of one sifting step in the construction of a IMF. (a) Input time series highlighting the local maxima and the local minima. (b) Time series with the interpolated upper and lower envelopes (slashed lines). (c) Time series with the envelopes and the mean of both envelopes (red line). (d) First iteration of the sifting process. In this example, the extracted function does not satisfy the IMF's conditions and therefore another set of sifting processes must be applied (see Fig.2).

3.1. Cross-correlations across time-scales

The proposed time-scale-dependent correlation computes the Pearson correlation coefficients between two components IMF_i^X, IMF_j^Y, $i, j = 1, \ldots, n$ obtained from the decomposition of the time series $X(t)$ and $Y(t)$, respectively:

$$\rho_{i,j}^{XY} = \frac{1}{T} \sum_{t=1}^{T} \frac{\left(\mathrm{IMF}_i^X(t) - \overline{\mathrm{IMF}_i^X}\right)\left(\mathrm{IMF}_j^Y(t) - \overline{\mathrm{IMF}_j^Y}\right)}{\sigma_i^X \sigma_j^Y} \tag{2}$$

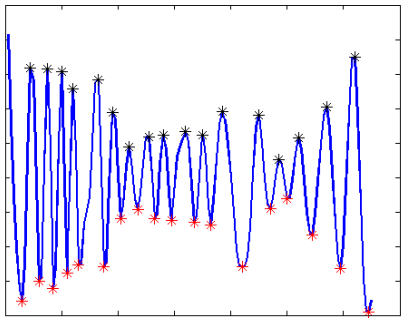

(a) Local maxima and minima.

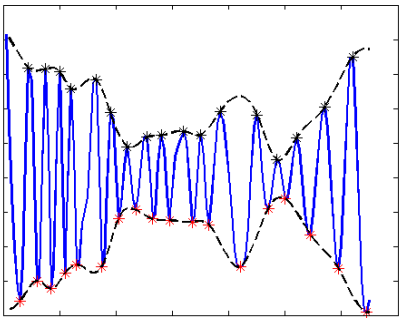

(b) Upper and lower envelopes.

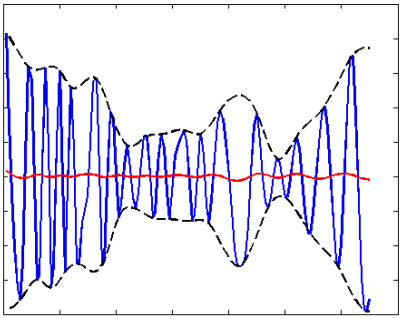

(c) Envelope mean.

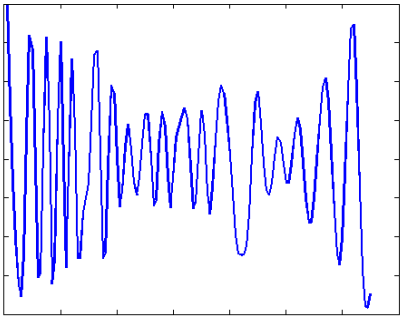

(d) IMF example.

Figure 2: Example of the few final sifting steps which produce a valid IMF. (a) Input time series highlighting the local maxima and the local minima. (b) Input time series with the interpolated upper and lower envelopes (slash lines). (c) Input time series with the envelopes and the mean of both envelopes (red line). (d) Last iteration of the sifting process, the extracted function is the first IMF.

where $\overline{\text{IMF}_i^X}$ denotes the sample mean over time of IMF_i^X and σ_i^X denotes the sample standard deviation of IMF_i^X.

Although the IMFs are not theoretically stationary, the IMFs satisfy the condition of having local mean equal to zero and can then be considered to be at least locally stationary [32]. Contrary, the residue does not need to satisfy the IMF conditions, and particularly, if we decompose a non-stationary time series, the extracted residue will contain the trend of the time series, making it a non-stationary component. Thus, a correlation coefficient between the residues is just a measure of linear dependency of the trends indicating

their direction of co-movement. This correlation coefficient is likely to be high and must not be directly compared with the correlations between the IMFs. Nonetheless, it is a significant part of the analysis, quantifying relative directions of the trends at the length scale of the time window.

3.2. Time-dependent lagged cross-correlation at the same time-scale over a rolling window

We also compute lagged cross-correlations over a rolling window, which for simplicity, we limit to the same time scale. The cross-correlations between two different time series $X(t)$ and $Y(t)$, lagged by λ, over rolling windows of size W and at the same time-scale component i is defined as:

$$\rho_i^{XY}(t,\lambda) = \frac{1}{W-\lambda} \sum_{\tau=t-W+1}^{t-\lambda} \frac{\left(\mathrm{IMF}_i^X(\tau) - \overline{\mathrm{IMF}_i^X}\right)\left(\mathrm{IMF}_i^Y(\tau+\lambda) - \overline{\mathrm{IMF}_i^Y}\right)}{\sigma_i^X \sigma_i^Y} \,. \tag{3}$$

The time-lag λ is measured in units of the sampling time-scale. The window approach has the advantage of only assuming local stationarity rather than stationarity over the entire time series. Although this method is based on a simple measure of correlation (Pearson correlation), it adapts to the nature of the data and provides a dynamic measure of correlation across time-scales.

4. Correlation analysis of intraday financial data

We consider intraday data sampled at 30-seconds intervals for two stock market indices and a volatility index, namely, the S&P 500 index (USA), the IPC index (Mexico) and the VIX index (implied volatility index, calculated by the Chicago Board Options Exchange, USA). The data was obtained from Bloomberg [36]. The observation period includes 184 days, ranging from September 2013 to July 2014 and it only considers the trading days available for all the three indices. Each day has 780 data points (6.5 hours). The effect of the sampling frequency was explored during a preliminary analysis and then settled to 30s as a reasonable example of high frequency data. We did not observe qualitative differences in the results when varying the frequency from a few seconds to a few minutes. Note that the S&P 500 and the VIX indices are closely related, being one a function of the other. Differently, the IPC index is related to a separated economic system. We expect to observe significant differences in the relative correlation structures across time scales for this different indices.

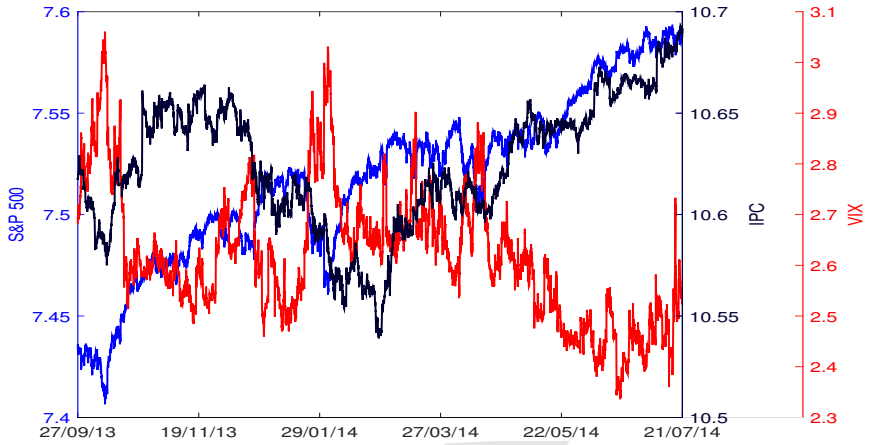

Figure 3: Intraday observations (sampled at 30-seconds intervals) for the S&P 500, the IPC and the VIX indices for the time period September 2013 to July 2014.

Figure 3 reports the dynamics of these three indices over that time period. We can observe that the S&P 500 and the IPC indices have similar behaviours. They are indeed positively correlated with correlation coefficient between log-returns equal to 0.21; this is in agreement with previous studies [37, 38]. On the contrary, the risk-price relationship between the S&P 500 and the VIX indices shows negative correlation, as reported for example in [39]. The correlation coefficient between these log-returns is equal to -0.26. Finally, the IPC and the VIX indices are essentially uncorrelated with very small negative correlations (the correlation coefficient between log-returns equals to -0.02).

4.1. Intraday analysis of correlation, example for the day July 18^{th} 2014

In order to easily visualize the results, let us exemplify the intraday analysis of the correlation on a randomly chosen day: July 18^{th} 2014. This day was chosen randomly and it appears as a very ordinary day, in this respect representative of many others. However, let us stress that we also performed the analysis for the complete dataset consisting of 184 days. The results are reported later in Section 4.2. Figure 4 displays the logarithm of prices for the three indices. Applying the EMD to each time series, we obtained five IMFs and a residue which are reported in Figure 5. The oscillating period of each IMFs is calculated by dividing the total number of points by the number of peaks, with rounded values reported in Table 1.

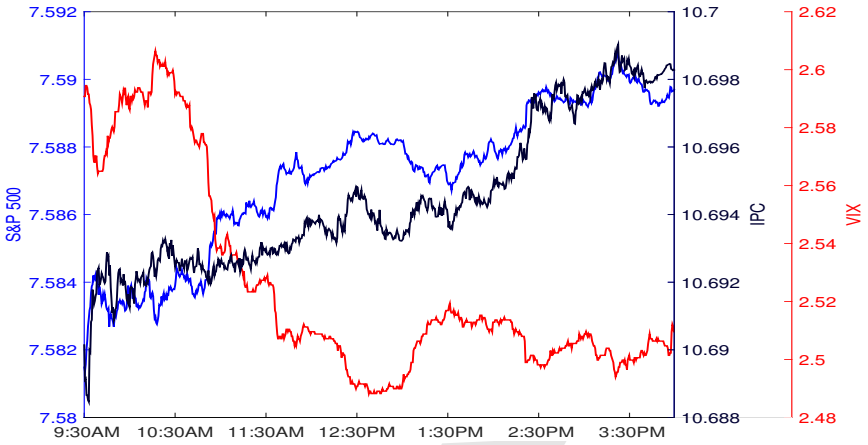

Figure 4: Intraday log-prices for the S&P 500, the IPC and the VIX indices, example for July 18^{th} 2014.

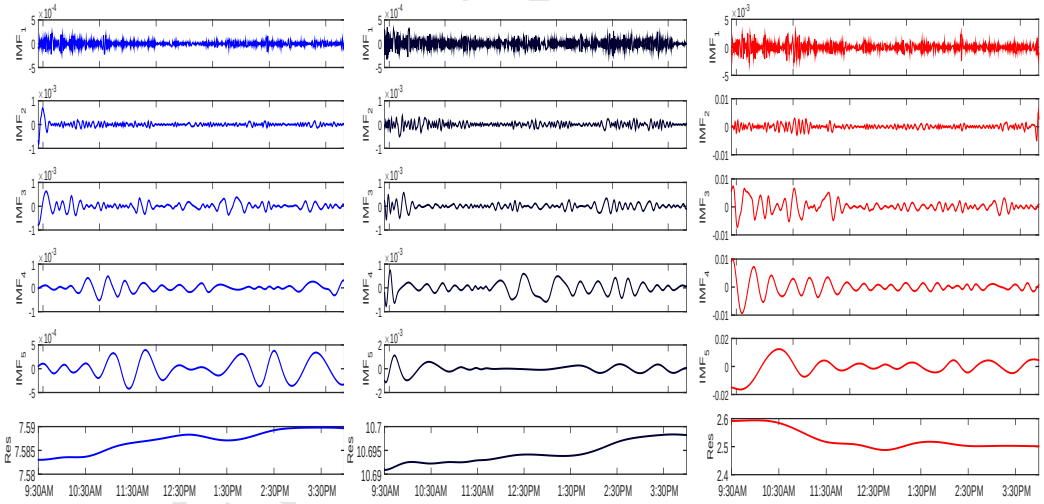

(a) S&P 500 index. (b) IPC index. (c) VIX index.

Figure 5: IMFs of the stock market indices and the volatility index, example for July 18^{th} 2014. From top to bottom IMF1...IMF4 and residue. X-axes time from 9:30 to 16:00.

4.1.1. Time-scale-dependent correlation, example for July 18^{th} 2014

We computed the time-scale-dependent correlation by means of Equation (2). The results are represented as a matrix of pairwise correlations between the IMFs where the magnitude of the correlation is visually represented by a color-map. Figure 6(a) shows the correlation matrix between the S&P 500

Index	IMF_1	IMF_2	IMF_3	IMF_4	IMF_5	**Residue**
S&P	4	8	20	44	88	–
IPC	4	8	16	40	88	–
VIX	4	8	20	40	88	–

Table 1: Oscillating period for the IMFs shown in Figure 5 and estimated by dividing the total number of points by the number of peaks and rounding, example for July 18^{th} 2014.

and the IPC indices. We observe positive correlations with mostly larger values on the diagonal (same time-scale components indices) with an increasing magnitude for increasing IMF time-scale. The correlation between the S&P and the VIX indices (Figure 6(b)) also reveals with larger values on the diagonal but instead with negative values between the residues and between the long time horizon components up to 8 minutes. Conversely, IMF_2 (4 minutes time-horizon, see Table 1) and IMF_1 (2 minutes time horizon) reveal instead small positive correlations.

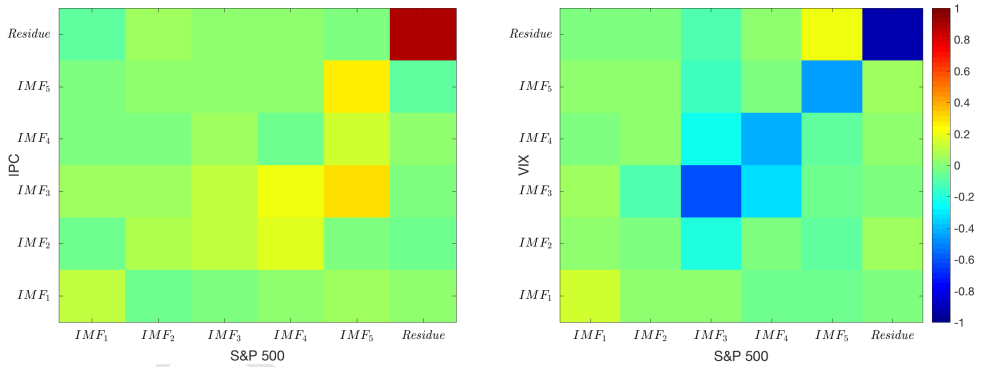

(a) S&P 500 index versus IPC index. (b) S&P 500 index versus VIX index.

Figure 6: Time-scale-dependent correlation structures, example for July 18^{th} 2014.

4.1.2. Time-dependent correlation, example for July 18^{th} *2014*

We also estimated the time-dependent lagged correlation between the S&P 500 and the IPC indices by using Equation (3). These correlations are represented in Figure 7(a) as a color-map matrix in which each column represents a successive window and each row represents a specific time-lag. The intraday correlation values are reported after W observations, with W

the size of the rolling-window. For the highest frequency components the rolling window size is 10 minutes whereas for the lowest frequency the window reaches one hour (see Table 2).

Lags are limited to $\lambda \leq \max(P_{X_i}, P_{Y_i})$, with P_{X_i} and P_{Y_i} denoting the oscillating period of IMF_i^X and IMF_i^Y, respectively. Choosing λ larger than the oscillating period results in repetitive patterns in the correlation structure. On the other hand, a shorter time-lag may not reveal some correlations. The window size is set at $W = \max(\lambda, 20)$.

From Figure 7(a), it is difficult to identify correlations patterns for the highest time-scale IMFs. However, for IMFs with lower time-scale, IMF_2, ..., IMF_5, we observe intervals of stronger correlations characterized by the nature of the oscillating IMFs, i.e., we observe lapses of positive correlation lagged in time by negative values of correlation, making the lead-lag relation between the IMFs almost symmetric with respect to the zero lag.

Figure 7(b) shows the correlation matrices for the S&P 500 and the VIX indices. Contrary to the correlation between the S&P 500 and the IPC indices, the correlation between the S&P and the VIX indices is negative at all frequencies and during the entire trading day. At the highest time-scale, IMF_1, we observe a clear pattern of negative correlation at lag $\lambda = 2$ (1 min), indicating that the S&P 500 leads the VIX index by 1 minute. When correlating the residue components, we observe a dominant blue band, indicating a negative correlation region (a similar red band is observed for the correlation between the S&P 500 and the IPC indices). Such a band could be attributed to the linear and non-stationary characteristics of the residues.

4.2. Intraday correlation, analysis on the complete dataset

Proceeding in the same way as illustrated in the previous examples for the particular day July 18^{th} 2014, we decomposed each daily time series into five IMFs and a residue. We then computed the time-scale-dependent correlation and the time-dependent correlation for each of the 184 days available in the dataset.

4.2.1. Time-scale-dependent correlation

The statistics for the time-scale dependent correlation between the IMFs with the same time-scale index for all trading days are reported in histograms. In Figure 8, we report histograms for the S&P 500 and the IPC indices and in Figure 9 for the S&P and the VIX indices. We observe prevalently positive

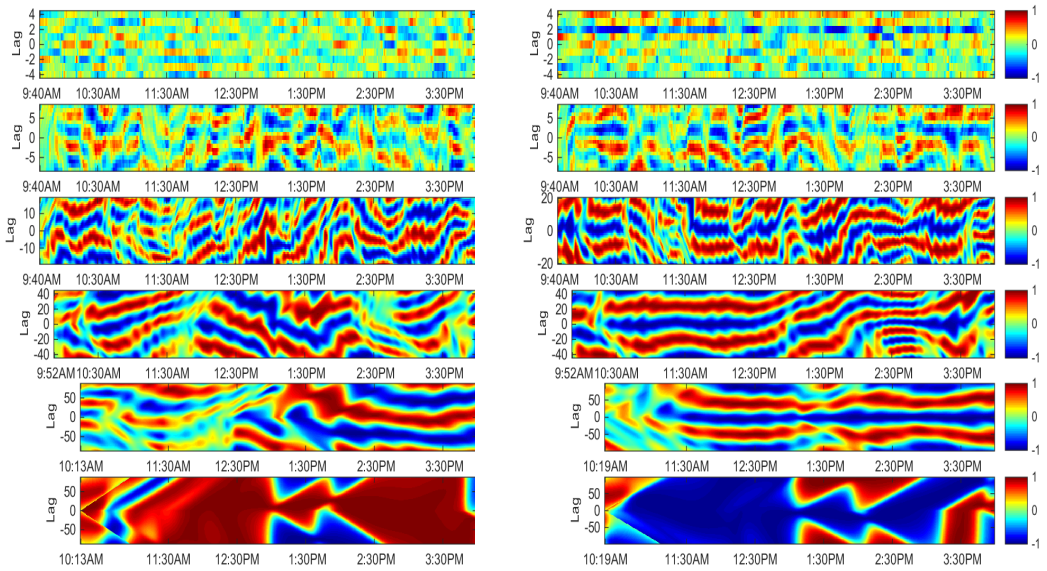

(a) S&P 500 index versus IPC index. (b) S&P 500 index versus VIX index.

Figure 7: Intraday time-dependent correlation, example for July 18^{th} 2014. (Rolling window sizes reported in in Table.2.)

correlations for the S&P 500 and the IPC components and instead prevalently negative correlations for the S&P 500 and the VIX. The histograms reveal significant deviations from zero for all the components with larger positive or negative correlations for components with longer time-scales.

Correlations between the IMF components with different time-scales indices are reported in Figure 10 where the sample median correlations are also reported. We use the sample median of the distribution since this statistic is not influenced by outliers. The case S&P and IPC is shown in Figure 10(a) and the S&P and VIX in Figure 10(b). The color-map matrices at the top are the median correlations between different time-scales whereas the plots below are the values of the diagonal elements (components with same-time scale indices).

4.2.2. Rolling window analysis and lag relations

We analysed the median of the time-dependent correlation matrices, computed as reported in Equation 3, over different rolling windows and various

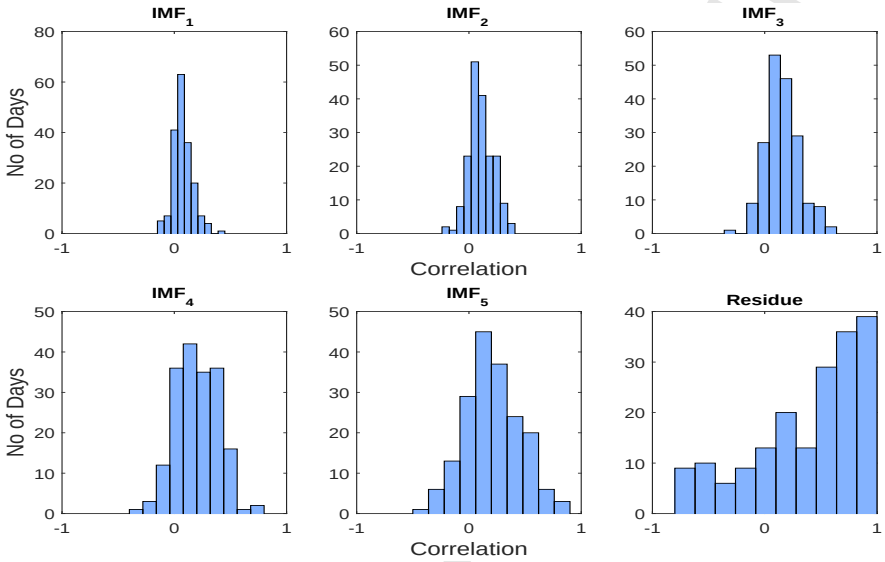

Figure 8: Distribution of the time-scale-dependent correlation between the IMFs of the S&P 500 index and the IMFs of the IPC index.

lags. The window sizes and the time lags which were used are reported in Table 2.

Component	**S&P vs IPC**		**S&P vs VIX**	
	Lag	**Window**	**Lag**	**Window**
IMF_1	4	20	4	20
IMF_2	9	20	9	20
IMF_3	19	20	21	21
IMF_4	44	44	48	48
IMF_5	110	110	124	124
Residue	110	110	124	124

Table 2: Average of the number of lags and the size of the rolling-window used for the time-dependent correlation analysis. Numbers are in 30 seconds units.

The median, time-varying, lagged correlation matrix (Eq.3) between the S&P 500 and the IPC indices is displayed in Figure 11(a). Overall, we observe relatively small correlations with little lead-lag relationships at all time-scales with larger values for the last two components and the residue (bottom panels). We observe patterns in the intraday activity with less persistent correlations around the middle of the day.

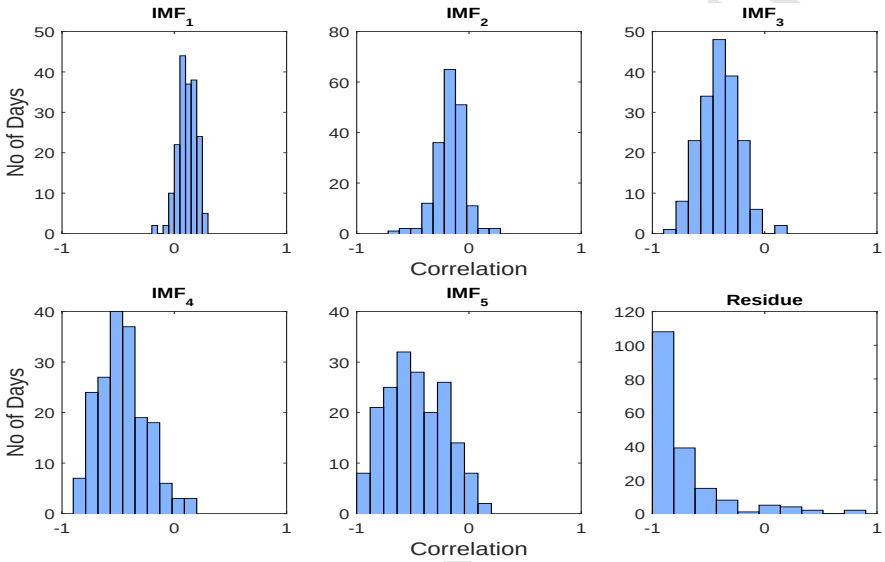

Figure 9: Distribution of the time-scale-dependent correlation between the IMFs of the S&P 500 index and the IMFs of the VIX index.

More intense negative correlation is observed between the S&P 500 and the VIX indices, reported in Figure 11(b). Interestingly, in this case, we observe significant lagged correlations at small time scales (IMF_1, IMF_2 and IMF_3) with the S&P 500 leading the VIX at one minute lag ($\lambda = 2$) with a stable pattern across the day. This indicates that consistently changes in the S&P 500 are followed by changes in the VIX after about 1 min and in the opposite direction (negative correlations).

5. Discussions and conclusions

In this paper we propose a simple approach which demonstrates that the Empirical Mode Decomposition can be used to investigate the correlation between time series at different time-scales. This expands the concept of correlations to a higher-dimensional level. We observed that, although most of the correlation is between components of the same time-scales, there are some significant correlations also between components of different time-scales. A dynamical analysis performed over rolling windows shows that correlations' patterns are both time and time-scale dependent. We uncovered lead-lag relations within components with the discovery of a persistent and significant 1 min negative coupling between the S&P 500 and the VIX

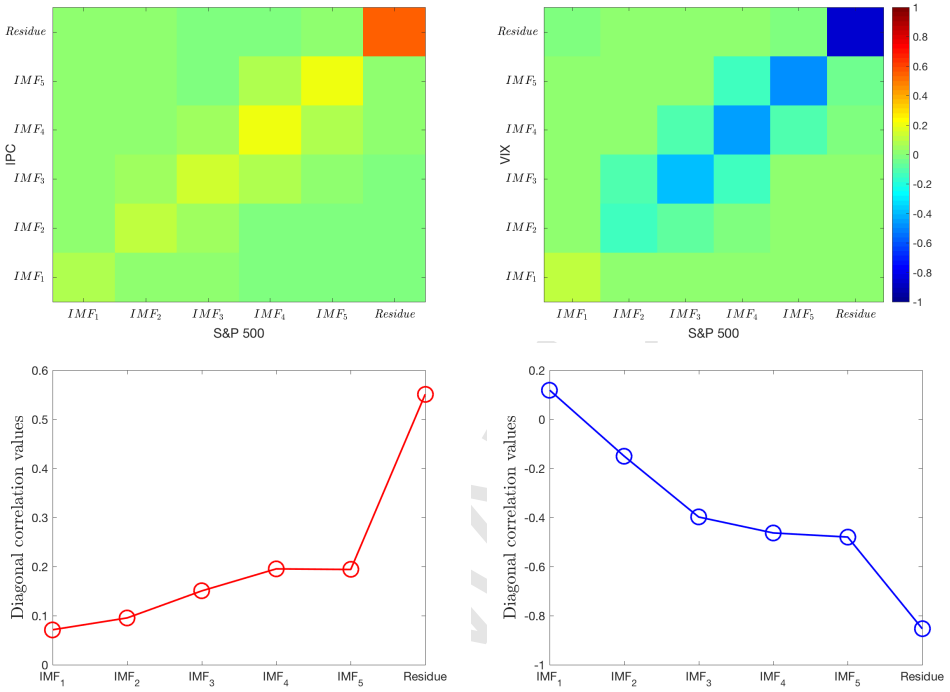

(a) S&P 500 index and the IPC index. (b) S&P 500 index and the VIX index.

Figure 10: Sample median of the time-scale-dependent correlation matrices over the time period from September 2013 to July 2014. The color-map matrices above are the cross correlations between components at different time-scales; the plots below report the values of the elements in the diagonal with same time-scales indices.

indices, which can have practical relevance for trading strategies and risk modeling.

In this paper we have proposed a simple methodology avoiding assumptions in the derivation of the intrinsic components and in the quantification of the correlations. The methodology we introduced in this paper and our findings are consistent with the coherence measure obtained with the wavelet transform [25]. However, given the simplicity of our approach and its adaptability to different time series without needing to specify any a priori filter function, we believe that the proposed correlation measures offer a simpler, computationally more efficient and easier to interpret approach.

The measures proposed in this paper and in particular Eqs. 2, 3 are the simplest generalisations of the linear correlation measure to include time-scale

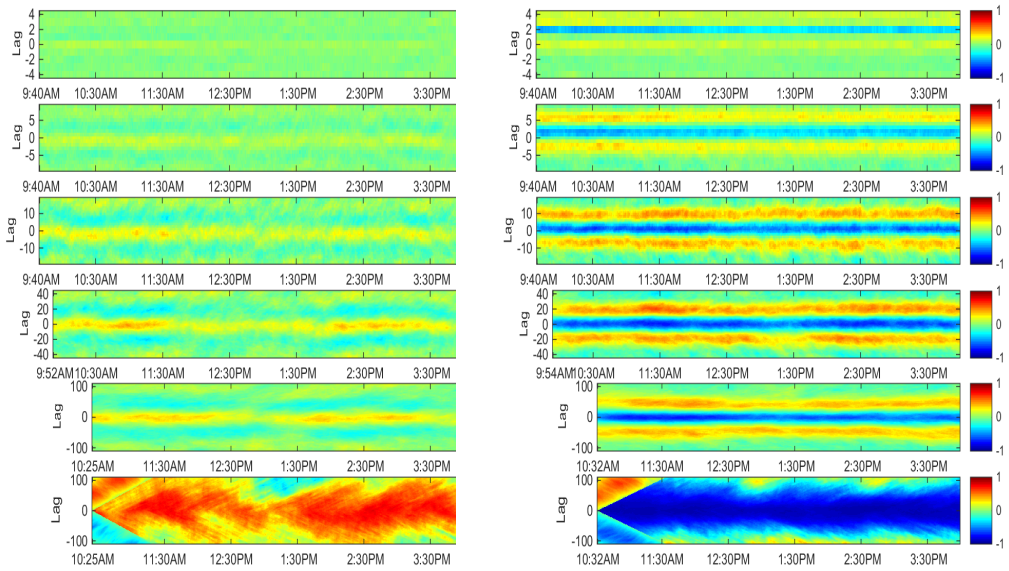

(a) S&P 500 index versus IPC index. (b) S&P 500 index versus VIX index.

Figure 11: Sample median of the time-dependent correlation matrices over the time period from September 2013 to July 2014. (Rolling window sizes reported in in Table.2.)

components. We chose them as a natural extension of the cross-correlation concept. However, there are some aspects of the present approach that would be interesting to further investigate in future works. For instance, Eqs. 2, 3 perform averages over the variables, but time-series with different oscillation scales lead to different averages even if the (scaled) nature of the variable is the same. This is probably penalising the values associated with high frequency components that we indeed observe to be consistently smaller. Similarly, the Pearson correlations between the residues must be interpreted with caution because -by construction- the residues are non-stationary and their correlation coefficient is a measure of the relative trend rather than a quantification of dependency.

Acknowledgement

The authors wish to thank Bloomberg for providing the data. NN would like to acknowledge the financial support from Conacyt-Mexico. TDM wishes

to thank the COST Action TD1210 for partially supporting this work. TA & TDM wish to thank the Systemic Risk Centre at LSE.

[1] Rosario N Mantegna and H Eugene Stanley. Introduction to econophysics: correlations and complexity in finance. Cambridge university press, 1999.

[2] Tomaso Aste, W Shaw, and T. Di Matteo. Correlation structure and dynamics in volatile markets. New Journal of Physics, 12(8):085009, 2010.

[3] Tomaso Aste and Tiziana Di Matteo. Dynamical networks from correlations. Physica A: Statistical Mechanics and its Applications, 370(1):156–161, 2006.

[4] Tomaso Aste and Tiziana Di Matteo. Introduction to complex and econophysics systems: A navigation map. Complex physical, biophysical and econophysical systems, pages 1–35, 2010.

[5] Rosario N Mantegna and H Eugene Stanley. An introduction to econophysics: correlation and complexity in finance, 2000.

[6] Michele Tumminello, Tiziana Di Matteo, Tomaso Aste, and Rosario N. Mantegna. Correlation based networks of equity returns sampled at different time horizons. The European Physical Journal B, 55(2):209–217, 2007.

[7] Christian Borghesi, Matteo Marsili, and Salvatore Micciche. Emergence of time-horizon invariant correlation structure in financial returns by subtraction of the market mode. Physical Review E, 76(2):026104, 2007.

[8] Daniel J Fenn, Mason A Porter, Stacy Williams, Mark McDonald, Neil F Johnson, and Nick S Jones. Temporal evolution of financial-market correlations. Physical review E, 84(2):026109, 2011.

[9] Chester Curme, Michele Tumminello, Rosario N. Mantegna, H. Eugene Stanley, and Dror Y. Kenett. Emergence of statistically validated financial intraday lead-lag relationships. Quantitative Finance, 15(8):1375–1386, 2015.

[10] Lisa Uechi, Tatsuya Akutsu, H Eugene Stanley, Alan J Marcus, and Dror Y Kenett. Sector dominance ratio analysis of financial markets. Physica A: Statistical Mechanics and its Applications, 421:488–509, 2015.

[11] Dror Y Kenett, Xuqing Huang, Irena Vodenska, Shlomo Havlin, and H Eugene Stanley. Partial correlation analysis: Applications for financial markets. Quantitative Finance, 15(4):569–578, 2015.

[12] Nicolo Musmeci, Tomaso Aste, and Tiziana Di Matteo. Relation between financial market structure and the real economy: Comparison between clustering methods. PloS one, 10(3):e0116201, 2015.

[13] Nicoló Musmeci, Vincenzo Nicosia, Tomaso Aste, Tiziana Di Matteo, and Vito Latora. The multiplex dependency structure of financial markets. Complexity, 2017, 2017.

[14] Tomaso Aste and T Di Matteo. Sparse causality network retrieval from short time series. Complexity, 2017, 2017.

[15] Tiziana Di Matteo. Multi-scaling in finance. Quantitative finance, 7(1):21–36, 2007.

[16] RJ Buonocore, T Aste, and T Di Matteo. Asymptotic scaling properties and estimation of the generalized hurst exponents in financial data. Physical Review E, 95(4):042311, 2017.

[17] RJ Buonocore, N Musmeci, T Aste, and T Di Matteo. Two different flavours of complexity in financial data. The European Physical Journal Special Topics, 225(17-18):3105–3113, 2016.

[18] Nicoló Musmeci, Tomaso Aste, and Tiziana Di Matteo. Risk diversification: a study of persistence with a filtered correlation-network approach. arXiv preprint arXiv:1410.5621, 2014.

[19] Nicoló Musmeci, Tomaso Aste, and T Di Matteo. Interplay between past market correlation structure changes and future volatility outbursts. Scientific reports, 6, 2016.

[20] Christoph Aymanns, Fabio Caccioli, J Doyne Farmer, and Vincent WC Tan. Taming the basel leverage cycle. Journal of Financial Stability, 27:263–277, 2016.

[21] Istvan Varga-Haszonits, Fabio Caccioli, and Imre Kondor. Replica approach to mean-variance portfolio optimization. Journal of Statistical Mechanics: Theory and Experiment, 2016(12):123404, 2016.

[22] Fabio Caccioli, Imre Kondor, Matteo Marsili, and Susanne Still. Liquidity risk and instabilities in portfolio optimization. International Journal of Theoretical and Applied Finance, 19(05):1650035, 2016.

[23] François Longin and Bruno Solnik. Is the correlation in international equity returns constant: 1960-1990? Journal of International Money and Finance, 14(1):3–26, 1995.

[24] Antonio Rua and Lucas C. Nunes. International comovement of stock market returns: A wavelet analysis. Journal of Empirical Finance, 16(4):632–639, 2009.

[25] Lukas Vácha and Jozef Baruník. Co-movement of energy commodities revisited: Evidence from wavelet coherence analysis. Energy Economics, 34(1):241–247, 2012.

[26] Marco Bartolozzi, Christopher Mellen, Tiziana Di Matteo, and Tomaso Aste. Multi-scale correlations in different futures markets. The European Physical Journal B, 58(2):207–220, 2007.

[27] Jean-François Muzy, Emmanuel Bacry, and Alain Arneodo. Multifractal formalism for fractal signals: The structure-function approach versus the wavelet-transform modulus-maxima method. Physical review E, 47(2):875, 1993.

[28] Levent Sendur and Ivan W Selesnick. Bivariate shrinkage functions for wavelet-based denoising exploiting interscale dependency. IEEE Transactions on signal processing, 50(11):2744–2756, 2002.

[29] Jan W Kantelhardt, Stephan A Zschiegner, Eva Koscielny-Bunde, Shlomo Havlin, Armin Bunde, and H Eugene Stanley. Multifractal detrended fluctuation analysis of nonstationary time series. Physica A: Statistical Mechanics and its Applications, 316(1):87–114, 2002.

[30] Davor Horvatic, H Eugene Stanley, and Boris Podobnik. Detrended cross-correlation analysis for non-stationary time series with periodic trends. EPL (Europhysics Letters), 94(1):18007, 2011.

[31] Xiaonei Zhang, Ming Zeng, and Qinghao Meng. Multivariate multifractal detrended fluctuation analysis of 3d wind field signals. Physica A: Statistical Mechanics and its Applications, 490:513–523, 2018.

[32] Norden E. Huang, Zheng Shen, Steven R. Long, Manli C. Wu, Hsing H. Shih, Quanan Zheng, Nai-Chyuan Yen, Chi Chao Tung, and Henry H. Liu. The empirical mode decomposition and the Hilbert spectrum for non-linear and non-stationary time series analysis. Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences, 454(1971):903–995, 1998.

[33] Z.K. Peng, Peter W. Tse, and F.L. Chu. A comparison study of improved Hilbert-Huang transform and wavelet transform: Application to fault diagnosis for rolling bearing. Mechanical Systems and Signal Processing, 19(5):974–988, 2005.

[34] X. Chen, Z. Wu, and N.E. Huang. The time-dependent intrinsic correlation based on the empirical mode decomposition. Advances in Adaptive Data Analysis, 2(2):233–265, 2010. cited By 19.

[35] P. Flandrin, G. Rilling, and P. Goncalves. Empirical mode decomposition as a filter bank. Signal Processing Letters, IEEE, 11(2):112–114, Feb 2004.

[36] https://www.bloomberg.com/notices/financial-data/.

[37] Eurilton Araújo. Macroeconomic shocks and the co-movement of stock returns in Latin America. Emerging Markets Review, 10(4):331–344, 2009.

[38] Mohamed El Héidi Arouri, Amine Lahiani, and Duc Khuong Nguyen. Cross-market dynamics and optimal portfolio strategies in Latin American equity markets. European Business Review, 27(2):161–181, 2015.

[39] Robert E Whaley. The investor fear gauge. The Journal of Portfolio Management, 26(3):12–17, 2000.