MDPI

Article

Local Gaussian Cross-Spectrum Analysis

Lars Arne Jordanger [1,*,†] and Dag Tjøstheim [2,†]

1 Department of Computer Science, Electrical Engineering and Mathematical Sciences, Faculty of Engineering and Science, Western Norway University of Applied Sciences, 5020 Bergen, Norway
2 Department of Mathematics, Faculty of Mathematics and Natural Sciences, University of Bergen, 5020 Bergen, Norway
* Correspondence: lars.arne.jordanger@hvl.no
† These authors contributed equally to this work.

Abstract: The ordinary spectrum is restricted in its applications, since it is based on the second-order moments (auto- and cross-covariances). Alternative approaches to spectrum analysis have been investigated based on other measures of dependence. One such approach was developed for univariate time series by the authors of this paper using the *local Gaussian auto-spectrum* based on the *local Gaussian auto-correlations*. This makes it possible to detect local structures in univariate time series that look similar to white noise when investigated by the ordinary auto-spectrum. In this paper, the *local Gaussian approach* is extended to a *local Gaussian cross-spectrum* for multivariate time series. The local Gaussian cross-spectrum has the desirable property that it coincides with the ordinary cross-spectrum for Gaussian time series, which implies that it can be used to detect non-Gaussian traits in the time series under investigation. In particular, if the ordinary spectrum is flat, then peaks and troughs of *the local Gaussian spectrum* can indicate nonlinear traits, which potentially might reveal *local periodic phenomena* that are undetected in an ordinary spectral analysis.

1. Introduction

The auto- and cross-spectra are tools that can extract temporal information from a time series, and they can be used to detect interesting periodic phenomena from the data under investigation. A disadvantage of these tools is that they are based on the auto- and cross-correlations. This implies that nonlinear temporal dependencies could be invisible from the point of view of the classical spectral theory, and this has, as discussed below, motivated quite a few modifications of the classical approach.

One such modification, the local Gaussian spectral approach, was introduced for the univariate case in Jordanger and Tjøstheim (2022) (hereafter referred to as JT22). In contradistinction to the ordinary spectrum, which is based on linear operations on the second moments and with a resulting estimate constructed from the periodogram, the locally Gaussian spectrum is distributional-based using local maximum likelihood estimation. The periodogram cannot be used in this local Gaussian framework. Nevertheless, a key feature of the local Gaussian approach is that the resulting local Gaussian spectra coincide completely with the classical spectra for Gaussian time series. The local Gaussian auto-spectra can be used to investigate how the temporal dependency structure for a given univariate time series deviates from the linear Gaussian dependency structure, and a local Gaussian analysis can thus detect nonlinear temporal dependencies, including local periodic phenomena, that are hidden from the classical approach.

This paper presents the extension of the local Gaussian approach to the multivariate case, introducing local Gaussian variants of the classical (global) co-spectrum, the quadrature-spectrum, the amplitude-spectrum, and the phase-spectrum.[1] Using these new

local Gaussian tools, it is possible to investigate, on a deeper than linear dependency level, how different time series are interconnected. The local Gaussian approach can detect cross-temporal interdependency relationships beyond what the classical (global) cross-spectrum can manage, and this can, e.g., be used to examine whether one series is leading–lagging another one at extreme amplitudes, but not, say, at moderate amplitudes.

The local Gaussian spectral-estimation algorithm in this paper was sanity tested on simulated examples (a nonlinearly connected sinusoidal example, as explained later on), and it was also used on the French CAC index versus the German DAX index in order to show how it can be applied to investigate nonlinear cross-temporal interdependency structures from the `EuStockMarkets`-data.

In addition to this, a GARCH-type model was fitted to the `EuStockMarkets`-data, and samples from the fitted model were used to show how the local Gaussian sanity testing of parametric models from JT22 can be extended to the multivariate case. This local Gaussian sanity testing of the fitted models could supplement other model selection techniques, and it could help detect models whose local Gaussian properties clearly deviate from the properties seen in the sample. Note that the focus for this particular part, as in JT22, is on the concepts themselves. A precise analysis of a test of fit is left to future work.

The classical (global) approach: The auto- and cross-covariances $\left\{\left\{\gamma_{k\ell}(h)\right\}_{h\in\mathbb{Z}}\right\}_{k,\ell=1}^{d}$ from a time series $\left\{\mathbf{Y}_t = \left(Y_{1,t}, \ldots, Y_{d,t}\right)\right\}_{t\in\mathbb{Z}}$, can range from determining it completely (Gaussian time series) to containing no information at all (GARCH-type models, where squaring of observations can be used as an ad hoc device to treat dependence). The auto- and cross-spectral densities $\left\{f_{k\ell}(\omega)\right\}_{k,\ell=1}^{d}$ are (assuming they exist) the Fourier transformations of the auto- and cross-covariances, and these tools thus share the same limitations. This implies that the auto- and cross-spectra might be inadequate tools when the task of interest is to investigate non-Gaussian time series containing asymmetries or other nonlinear structures—such as those observed in stock returns, cf., e.g., Hong et al. (2007).

In ordinary spectral analysis, if $\left\{Y_{k,t}\right\}_{t\in\mathbb{Z}}$ and $\left\{Y_{\ell,t}\right\}_{t\in\mathbb{Z}}$ are jointly weakly stationary, and if the cross-covariances $\gamma_{k\ell}(h) := \mathrm{Cov}\left(Y_{k,t+h}, Y_{\ell,t}\right)$ are absolutely summable, then the cross-spectrum $f_{k\ell}(\omega)$ is defined as the Fourier transform of the cross-covariances, i.e.,

$$f_{k\ell}(\omega) = \sum_{h\in\mathbb{Z}} \gamma_{k\ell}(h) \cdot e^{-2\pi i\omega h}. \tag{1}$$

A simple scaling connects the auto- and cross-covariances $\gamma_{k\ell}(h)$ and the corresponding auto- and cross-correlations $\rho_{k\ell}(h)$, in particular $\gamma_{k\ell}(h) = \sqrt{\gamma_{kk}(0) \cdot \gamma_{\ell\ell}(0)} \cdot \rho_{k\ell}(h)$, and this implies that the cross-spectrum $f_{k\ell}(\omega)$ given in Equation (1) also can be written as

$$f_{k\ell}(\omega) = \sqrt{\gamma_{kk}(0) \cdot \gamma_{\ell\ell}(0)} \cdot \sum_{h\in\mathbb{Z}} \rho_{k\ell}(h) \cdot e^{-2\pi i\omega h}. \tag{2}$$

It is the sum of the cross-correlations that are generalized in the local Gaussian approach.

The expression for the inverse Fourier transform reveals, when $h = 0$, that the covariance $\mathrm{Cov}\left(Y_{k,t}, Y_{\ell,t}\right) = \gamma_{k\ell}(0)$ can be expressed as the integral $\int_{-1/2}^{1/2} f_{k\ell}(\omega)\, d\omega$. This makes it possible to inspect how the (linear) interaction between the marginal time series varies with the frequency ω. An inspection of the cross-spectrum $f_{k\ell}(\omega)$ is a bit more complicated than that of the auto-spectrum, since $f_{k\ell}(\omega)$, in general, will be a complex-valued function. It is thus usually the following real valued functions that are investigated:

$$c_{k\ell}(\omega) = \mathrm{Re}(f_{k\ell}(\omega)), \qquad\qquad q_{k\ell}(\omega) = -\,\mathrm{Im}(f_{k\ell}(\omega)), \tag{3a}$$
$$\alpha_{k\ell}(\omega) = \mathrm{Mod}(f_{k\ell}(\omega)), \qquad\qquad \phi_{k\ell}(\omega) = \mathrm{Arg}(f_{k\ell}(\omega)), \tag{3b}$$

where $c_{k\ell}(\omega)$, $q_{k\ell}(\omega)$, $\alpha_{k\ell}(\omega)$, and $\phi_{k\ell}(\omega)$, respectively, are referred to as the co-spectrum, quadrature-spectrum, amplitude-spectrum, and phase-spectrum. Note that $c_{k\ell}(\omega)$ always integrates to $\gamma_{k\ell}(0)$ over one period, whereas $q_{k\ell}(\omega)$ always integrates to zero.

The *coherence* $\mathcal{K}_{k\ell}(\omega) := f_{k\ell}(\omega)/\sqrt{f_{kk}(\omega)f_{\ell\ell}(\omega)}$ is an important tool when a spectral analysis is performed on a multivariate time series, in particular since $\mathcal{K}_{k\ell}(\omega)$ can be realized as the correlation of $\mathrm{d}Z_k(\omega)$ and $\mathrm{d}Z_\ell(\omega)$, where $Z_k(\omega)$ and $Z_\ell(\omega)$ are the right continuous orthogonal-increment processes that by the spectral representation theorem correspond to the weakly stationary time series $\left\{Y_{k,t}\right\}_{t\in\mathbb{Z}}$ and $\left\{Y_{\ell,t}\right\}_{t\in\mathbb{Z}}$ (see, e.g., Brockwell and Davis 1986, p. 436 for details). The *squared coherence* $\left|\mathcal{K}_{k\ell}(\omega)\right|^2$ is of interest since its value (in the interval $[0,1]$) reveals to what extent the two time series $\left\{Y_{k,t}\right\}_{t\in\mathbb{Z}}$ and $\left\{Y_{\ell,t}\right\}_{t\in\mathbb{Z}}$ can be related by a linear filter.

Modifications of the classical approach: Other spectral approaches, involving different generalizations of the auto-spectrum $f(\omega)$, were discussed in JT22 [Section 1], and several of the approaches were based on the following idea: The second-order moments captured by the autocovariances $\{\gamma(h)\}_{h\in\mathbb{Z}}$ can be replaced by alternative dependence measures ξ_h computed from the bivariate random variables (Y_{t+h}, Y_t), and a spectral density approach can then (under suitable regularity conditions) be defined as the Fourier transform of $\{\xi_h\}_{h\in\mathbb{Z}}$. For multivariate time series, the natural extension is then to define similar measures $\xi_{k\ell:h}$ for the bivariate random variables $\left(Y_{k,t+h}, Y_{\ell,t}\right)$, and then use the corresponding Fourier transform as an alternative to the cross-spectrum $f_{k\ell}(\omega)$.

It does not seem to be the case (yet) that multivariate versions have been investigated for all of the possible generalizations of the auto-spectrum $f(\omega)$, but some generalizations do exist. The first extension of the cross-spectrum $f_{k\ell}(\omega)$ along these lines is the *polyspectra* introduced in Brillinger (1965), which is the multivariate version of the higher-order moments/cumulants approach to spectral analysis (see Tukey 1959; Brillinger 1984, 1991). Another generalization of $f_{k\ell}(\omega)$ is given in Chung and Hong (2007), where the generalized function approach introduced in Hong (1999) is used to set up a cross-spectrum that can be used for the testing of directional predictability in foreign exchange markets. Some recent work applies copula techniques for the investigation of multivariate time series, such as Li (2023) using a *copula spectra density kernel* and Zhao et al. (2022) using *copula-linked univariate D-vines*. A quantile-based approach can be found in Baruník and Kley (2019), where the *quantile-coherency* concept is introduced, and quantile approaches based on the *quantile periodogram* can be seen in Li (2021, 2022a).

In many cases, it can be difficult to interpret the various spectra/visualizations which result from these modifications of the classical approach, but machine learning methods can be combined with these techniques and help reduce/remove this problem—see, e.g., Chen et al. (2021); Li (2020, 2022b) for some examples of this approach. The review paper Ciaburro and Iannace (2021) contains more information about machine-learning-based methods applied to time series data.

The local Gaussian approach: The local Gaussian spectral density $f_v(\omega)$ for univariate strictly stationary time series that was defined in JT22 is based on the local Gaussian auto-correlations $\rho_v(h)$ from Tjøstheim and Hufthammer (2013). A simple adjustment gives the local Gaussian cross-correlations $\rho_{k\ell:v}(h)$ for multivariate strictly stationary time series, from which a local Gaussian analogue $f_{k\ell:v}(\omega)$ of the cross-spectrum $f_{k\ell}(\omega)$ can be constructed using the Fourier transform. The local Gaussian version of the cross-spectrum enables local Gaussian alternatives to be defined of the co-spectrum, quadrature-spectrum, amplitude-spectrum, and phase-spectrum, by simply copying the setup used in the ordinary (global) case. Local Gaussian analogues of the coherence and squared coherence were investigated in the preparation for this paper, but then discarded (see the discussion at the end of Section 2.3 for further details).

It should be mentioned that the local Gaussian approach to statistical modeling has been applied recently to a number of nonlinear statistical problems; see, e.g., Tjøstheim et al. (2021).

An overview of the paper: Section 2 defines the *local Gaussian cross-spectrum* $f_{k\ell:v}(\omega)$, which immediately gives the related local Gaussian variants of the co-spectrum, quadrature-spectrum, amplitude-spectrum, and phase-spectrum from Equation (3). The asymptotic theory for the estimators are then presented (the technical details and proofs are relegated to the Supplementary Material). The real and simulated examples in Section 3 show that estimates of $f_{k\ell:v}(\omega)$ can be used to detect and investigate nonlinear structures in non-Gaussian white noise, and in particular that $f_{k\ell:v}(\omega)$ can detect local periodic phenomena that are undetected in an ordinary spectral analysis. Some concluding remarks are gathered in Section 4.

Supplementary Material: An online Supplementary Material exists that includes additional in-depth discussions. See the end of the manuscript for further details about the content of the Supplementary Material.

2. Definitions

2.1. The Local Gaussian Correlations

At the core of the generalization of Equation (1) lies the local Gaussian correlation ρ_v from Tjøstheim and Hufthammer (2013). A construction that was originally used for density estimation (see Hjort and Jones 1996) was, in Tjøstheim and Hufthammer (2013), adjusted a bit in order to define ρ_v as a new local measure of dependence for a bivariate random variable W. With $g(w)$ being the density function of W, the idea is to find a local Gaussian approximation $\psi(w; \theta)$ in the neighborhood of v—i.e., start with

$$\psi(w; \theta) := \frac{1}{2\pi \cdot \sigma_1 \sigma_2 \sqrt{1 - \rho^2}} \exp\left\{ -\frac{\sigma_2^2 (w_1 - \mu_1)^2 - 2\sigma_1 \sigma_2 \rho (w_1 - \mu_1)(w_2 - \mu_2) + \sigma_1^2 (w_2 - \mu_2)^2}{\sigma_1^2 \sigma_2^2 (1 - \rho^2)} \right\}, \quad (4)$$

then find a parameter vector $\theta_v = \left[\mu_{1v}, \mu_{2v}, \sigma_{1v}, \sigma_{2v}, \rho_v \right]'$ that gives the best match to $g(w)$ in a neighborhood of v, and then extract the correlation component ρ_v from the parameter vector θ_v.

As discussed in JT22 [Section 2.1.2], the theoretical treatment can be based directly on the construction from Tjøstheim and Hufthammer (2013), but the numerical convergence of the estimation algorithm might then sometimes fail. This estimation problem can be countered if a *normalization of the marginals* (see Definition 1 below) is performed before the estimation algorithm is used.

2.2. The Local Gaussian Cross-Spectrum

The definition of the local Gaussian cross-spectrum density is almost identical to the definition of the local Gaussian spectral density from JT22 [Section 2.2], which, in this paper, will henceforth be referred to as the local Gaussian auto-spectrum.

Definition 1. *For a strictly stationary multivariate time series $\{Y_t\}_{t \in \mathbb{Z}}$, where $Y_t = \left(Y_{1,t}, \ldots, Y_{d,t} \right)$, the local Gaussian cross-spectrum of the marginal time series $\{Y_{k,t}\}_{t \in \mathbb{Z}}$ and $\{Y_{\ell,t}\}_{t \in \mathbb{Z}}$ is constructed in the following manner.*

1. With G_k and G_ℓ as the univariate marginal cumulative distributions of, respectively, $\{Y_{k,t}\}_{t \in \mathbb{Z}}$ and $\{Y_{\ell,t}\}_{t \in \mathbb{Z}}$, and Φ as the cumulative distribution of the univariate standard normal distribution, define normalized versions $\{Z_{k,t}\}_{t \in \mathbb{Z}}$ and $\{Z_{\ell,t}\}_{t \in \mathbb{Z}}$ by

$$\left\{ Z_{k,t} := \Phi^{-1}\left(G_k\left(Y_{k,t} \right) \right) \right\}_{t \in \mathbb{Z}}, \qquad \left\{ Z_{\ell,t} := \Phi^{-1}\left(G_\ell\left(Y_{\ell,t} \right) \right) \right\}_{t \in \mathbb{Z}}. \quad (5)$$

2. For a given point $v = (v_1, v_2)$ and for each bivariate lag h pair $Z_{k\ell:h:t} := (Z_{k:t+h}, Z_{\ell:t})$, a local Gaussian cross-correlation $\rho_{k\ell:v}(h)$ can be computed based on a five-parameter local Gaussian approximation of the bivariate density of $Z_{k\ell:h:t}$ at (v_1, v_2).

3. When $\sum_{h \in \mathbb{Z}} |\rho_{k\ell:v}(h)| < \infty$, the local Gaussian cross-spectrum at the point v is defined as

$$f_{k\ell:v}(\omega) := \sum_{h=-\infty}^{\infty} \rho_{k\ell:v}(h) \cdot e^{-2\pi i \omega h}. \tag{6}$$

The normalization of the marginals in Definition 1 reduces the amount of numerical convergence issues in the estimation algorithm, and the resulting local Gaussian cross-spectrum $f_{k\ell:v}(\omega)$ is thus a feature of the collection of bivariate copula structures that are present inside the overall cross-temporal dependency structure of the strictly stationary time series $\{Y_t\}_{t \in \mathbb{Z}}$. This implies that investigations based on the local Gaussian cross-spectrum $f_{k\ell:v}(\omega)$ should be complemented with methods that can extract relevant features from the probability density functions of the d marginal distributions.

The basic properties of the local Gaussian cross-spectrum are quite similar to those encountered for the local Gaussian auto-spectrum in JT22 [Lemma 2.1].

Lemma 1. *The following properties hold for $f_{k\ell:v}(\omega)$.*

1. $f_{k\ell:v}(\omega)$ *coincides with $f_{k\ell}(\omega)$ for all $v \in \mathbb{R}^2$ when $\{Y_t\}_{t \in \mathbb{Z}}$ is a multivariate Gaussian time series.*

2. *The following holds when $\breve{v} := (v_2, v_1)$ is the diagonal reflection of $v = (v_1, v_2)$:*

$$f_{k\ell:v}(\omega) = \overline{f_{\ell k:\breve{v}}(\omega)}, \tag{7a}$$

$$f_{k\ell:v}(\omega) = \rho_{k\ell:v}(0) + \sum_{h=1}^{\infty} \rho_{\ell k:\breve{v}}(h) \cdot e^{+2\pi i \omega h} + \sum_{h=1}^{\infty} \rho_{k\ell:v}(h) \cdot e^{-2\pi i \omega h}. \tag{7b}$$

Proof. Item 1 follows since the local Gaussian cross-correlations $\rho_{k\ell:v}(h)$ by construction coincides with the ordinary (global) cross-correlations $\rho(h)$ in the Gaussian case. For the proof of Item 2, the key observation is that the *diagonal folding property* that was observed for the local Gaussian auto-spectrum (see JT22 [Lemma C.1]) extends directly to the present case, i.e., $\rho_{k\ell:v}(-h) = \rho_{\ell k:\breve{v}}(h)$, where $\breve{v} = (v_2, v_1)$ is the *diagonally reflected point corresponding to v*. This implies that $f_{k\ell:v}(\omega) = \overline{f_{k\ell:v}(-\omega)} = \overline{f_{\ell k:\breve{v}}(\omega)}$, and it also follows that Equation (6) can be re-expressed as Equation (7b). □

The *diagonal folding property* enables both the estimation algorithm (see Algorithm 1) and the theoretical analysis (see Section 2.5) to be expressed in terms of non-negative lag-values. The discussion in Sections 2.3 and 2.4, in particular Definitions 2 and 3 and Algorithm 1, will thus use the same setup as that seen in Equation (7b).

2.3. Related Local Gaussian Entities

From the definition of the local Gaussian cross-spectrum, it is possible to define related spectra in the same manner as those mentioned for the ordinary spectrum in Equation (3).

Definition 2. *The local Gaussian versions of the co-spectrum $c_{k\ell}(\omega)$, the quadrature-spectrum $q_{k\ell}(\omega)$, the amplitude-spectrum $\alpha_{k\ell}(\omega)$, and the phase-spectrum $\phi_{k\ell}(\omega)$ are given by*

$$c_{k\ell:v}(\omega) := \text{Re}(f_{k\ell:v}(\omega)) = \sum_{h=-\infty}^{\infty} \rho_{k\ell:v}(h) \cdot \cos(2\pi\omega h)$$

$$= \rho_{k\ell:v}(0) + \sum_{h=1}^{\infty} \cos(2\pi\omega h)[\rho_{k\ell:v}(h) + \rho_{k\ell:\breve{v}}(h)], \tag{8a}$$

$$q_{k\ell:v}(\omega) := -\text{Im}(f_{k\ell:v}(\omega)) = \sum_{h=-\infty}^{\infty} \rho_{k\ell:v}(h) \cdot \sin(2\pi\omega h)$$

$$= \sum_{h=1}^{\infty} \sin(2\pi\omega h)[\rho_{k\ell:v}(h) - \rho_{k\ell:\breve{v}}(h)], \tag{8b}$$

$$\alpha_{k\ell:v}(\omega) := \text{Mod}(f_{k\ell:v}(\omega)) = \sqrt{c_{k\ell:v}^2(\omega) + q_{k\ell:v}^2(\omega)}, \tag{8c}$$

$$\phi_{k\ell:v}(\omega) := \text{Arg}(f_{k\ell:v}(\omega)) \in (-\pi, \pi]. \tag{8d}$$

The sums in Equation (8a,b) follow from Equation (7b). Note that Equation (7a) implies the following symmetry results: $c_{k\ell:v}(\omega) = c_{\ell k:\breve{v}}(\omega)$, $q_{k\ell:v}(\omega) = -q_{\ell k:\breve{v}}(\omega)$, $\alpha_{k\ell:v}(\omega) = \alpha_{\ell k:\breve{v}}(\omega)$ and $\phi_{k\ell:v}(\omega) = -\phi_{\ell k:\breve{v}}(\omega)$.

Algorithm 1 *For a sample* $\left\{ \boldsymbol{y}_t = \left(y_{1,t}, \ldots, y_{d,t} \right) \right\}_{t=1}^{n}$ *of size n from a multivariate time series, an m-truncated estimate* $\hat{f}_{k\ell:v}^{m}(\omega)$ *of* $f_{k\ell:v}(\omega)$ *is constructed by means of the following procedure.*

1. *Use the univariate marginals* $\left\{ y_{k,t} \right\}_{t=1}^{n}$ *and* $\left\{ y_{\ell,t} \right\}_{t=1}^{n}$ *to find estimates* $\widehat{G}_{k:n}$ *and* $\widehat{G}_{\ell:n}$ *of the corresponding marginal cumulative distribution functions, and compute from this the pseudo-normalized observations* $\left\{ \widehat{z}_{k,t} := \Phi^{-1}\left(\widehat{G}_{k:n}\left(y_{k,t} \right) \right) \right\}_{t=1}^{n}$ *and* $\left\{ \widehat{z}_{\ell,t} := \Phi^{-1}\left(\widehat{G}_{\ell:n}\left(y_{\ell,t} \right) \right) \right\}_{t=1}^{n}$.

2. *Create the lag h pseudo-normalized pairs* $\left\{ \left(\widehat{z}_{k,t+h}, \widehat{z}_{\ell,t} \right) \right\}_{t=1}^{n-h}$ *for* $h = 0, \ldots, m$, *and estimate for the point* $\boldsymbol{v} = (v_1, v_2)$ *the local Gaussian cross-correlations* $\left\{ \hat{\rho}_{k\ell:v}(h|b_h) \right\}_{h=0}^{m}$, *where the* $\left\{ b_h \right\}_{h=0}^{m}$ *is the bandwidths that are used for the different lags.*

3. *Create the lag h pseudo-normalized pairs* $\left\{ \left(\widehat{z}_{\ell,t+h}, \widehat{z}_{k,t} \right) \right\}_{t=1}^{n-h}$ *for* $h = 1, \ldots, m$, *and estimate for the diagonally reflected point* $\breve{\boldsymbol{v}} = (v_2, v_1)$ *the local Gaussian cross-correlations* $\left\{ \hat{\rho}_{\ell k:\breve{v}}(h|b_h) \right\}_{h=0}^{m}$.

4. *Adjust Equation (7b) from Lemma 1(2) with some lag-window function* $\lambda_m(h)$ *to obtain the estimate*

$$\hat{f}_{k\ell:v}^{m}(\omega) := \hat{\rho}_{k\ell:v}(0) + \sum_{h=1}^{m} \lambda_m(h) \cdot \hat{\rho}_{\ell k:\breve{v}}(h) \cdot e^{+2\pi i\omega h} + \sum_{h=1}^{m} \lambda_m(h) \cdot \hat{\rho}_{k\ell:v}(h) \cdot e^{-2\pi i\omega h}, \tag{9}$$

where the $\left\{ b_h \right\}_{h=0}^{m}$ *is suppressed from the notation in order to obtain a more compact formula.*

For Gaussian distributions, the local Gaussian correlations will always be equal to the ordinary (global) correlations[2], and the local Gaussian constructions in Definitions 1 and 2 will thus coincide with the ordinary (global) versions for multivariate Gaussian time series. A comparison of the local and global estimates in the same plot is thus of interest when a given sample is considered, since this could detect nonlinear dependencies of the time series under investigation.

It is possible to define a local Gaussian analogue of the squared coherence mentioned in Section 1 by replacing the ordinary cross- and auto-spectra with the corresponding local Gaussian versions, i.e., the object of interest would be $\mathcal{Q}_{k\ell:v}(\omega) := f_{k\ell:v}(\omega)f_{\ell k:v}(\omega)/f_{kk:v}(\omega)f_{\ell\ell:v}(\omega)$. This approach was investigated in the preparation of this paper, but it is not included here since $\mathcal{Q}_{k\ell:v}(\omega)$, in general, lacked the nice properties known from the

ordinary global case. In particular, the local Gaussian auto-spectra $f_{kk:v}(\omega)$ and $f_{\ell\ell:v}(\omega)$ will, in general, be complex-valued functions, so an inspection of $\mathcal{Q}_{k\ell:v}(\omega)$ must thus be based on plots of its real and imaginary parts (or its amplitude and phase). Moreover, these plots more often than not turned out to be rather hard to investigate, since the estimates of $f_{kk:v}(\omega)$ and $f_{\ell\ell:v}(\omega)$ (for some distributions and some frequencies ω) gave values very close to zero in the denominator.

2.4. Estimation

The estimation of the local Gaussian cross-spectrum $f_{k\ell:v}(\omega)$ from Section 2.2 follows the same setup that was used in JT22 [Section 2.3] for the estimation of the local Gaussian auto-spectrum, but some extra indices are needed in the present case. The estimates of the related spectra $c_{k\ell:v}(\omega)$, $q_{k\ell:v}(\omega)$, $\alpha_{k\ell:v}(\omega)$ and $\phi_{k\ell:v}(\omega)$ from Section 2.3 are then obtained from the estimate of $f_{k\ell:v}(\omega)$ in the obvious manner.

Definition 3. *For a multivariate sample $\{\mathbf{y}_t\}_{t=1}^n$ of size n, as described in Algorithm 1, the m-truncated estimates of the local Gaussian versions of the co-spectrum, quadrature-spectrum, amplitude-spectrum, and phase-spectrum are given by*

$$\hat{c}_{k\ell:v}^m(\omega) := \mathrm{Re}\left(\hat{f}_{k\ell:v}^m(\omega)\right) = \hat{\rho}_{k\ell:v}(0) + \sum_{h=1}^m \cos(2\pi\omega h)\lambda_m(h)[\hat{\rho}_{k\ell:v}(h) + \hat{\rho}_{k\ell:\ddot{v}}(h)], \quad (10a)$$

$$\hat{q}_{k\ell:v}^m(\omega) := -\mathrm{Im}\left(\hat{f}_{k\ell:v}^m(\omega)\right) = \sum_{h=1}^m \sin(2\pi\omega h)\lambda_m(h)[\hat{\rho}_{k\ell:v}(h) - \hat{\rho}_{k\ell:\ddot{v}}(h)], \quad (10b)$$

$$\hat{\alpha}_{k\ell:v}^m(\omega) := \mathrm{Mod}\left(\hat{f}_{k\ell:v}^m(\omega)\right) = \sqrt{\left(\hat{c}_{k\ell:v}^m(\omega)\right)^2 + \left(\hat{q}_{k\ell:v}^m(\omega)\right)^2}, \quad (10c)$$

$$\hat{\phi}_{k\ell:v}^m(\omega) := \mathrm{Arg}\left(\hat{f}_{k\ell:v}^m(\omega)\right) \in (-\pi, \pi]. \quad (10d)$$

The comments in JT22 [Section 2.3] hold for the present case, too. In particular, the estimated marginal cumulative distributions $\widehat{G}_{k:n}$ and $\widehat{G}_{\ell:n}$ from Algorithm 1(1) can either be based on the (rescaled) empirical cumulative distribution functions or they could be built upon a logspline technique such as the one implemented in Otneim and Tjøstheim (2017). Furthermore, for the asymptotic investigation, the arguments in Otneim and Tjøstheim (2017, Section 3) reveal that the pseudo-normalization of the marginals does not affect the final convergence rates, which (as was performed in JT22) implies that the present theoretical analysis can ignore the distinction between the original observations and the pseudo-normalized observations.

The Input Parameters and Some Other Technical Details

An inspection of Algorithm 1 reveals that several tuning parameters must be selected in order to compute the *m*-truncated local Gaussian cross-spectrum density estimates $\hat{f}_{k\ell:v}^m(\omega)$, and it is thus of interest to briefly present the values used for the plots contained in this paper. All simulated time series have the same length as the one encountered for the real data example, i.e., the log-returns of the `EuStockMarkets`-data.

Note that these tuning parameters were selected in order to provide a *proof-of-concept* for the fact that nonlinear dependency structures can be detected by this approach, and the quest for "optimal parameters" is a topic for further work. The interested reader can consult Section S3 in the Supplementary Material for a sensitivity analysis of the different tuning parameters.

The pseudo-normalization: The initial step of the computation of $\hat{f}_{k\ell:v}^m(\omega)$ is to replace the observations $\left\{\left(y_{1,t}, \ldots, y_{d,t}\right)\right\}_{t=1}^n$ with the corresponding pseudo-normalized observations $\left\{\left(\hat{z}_{1,t}, \ldots, \hat{z}_{d,t}\right)\right\}_{t=1}^n$, cf. Algorithm 1, that is, estimates of the d marginal cumulative density functions $\{G_\ell\}_{\ell=1}^d$ are needed. The present analysis used the rescaled empirical cumulative density functions $\hat{G}_{\ell:n}$ for this purpose.

The points v of investigation: Three diagonal points, with coordinates corresponding to the 10%, 50%, and 90% percentiles of the standard normal distribution,[3] are used in the basic plots in this section. Refer to Section S3.3 for further details related to the selection of v, and see Figures 2, 4 and 7 for some heatmap-based plots.

The lag-window function $\lambda_m(h)$: The smoothing of the estimated local Gaussian auto-correlations, cf. Algorithm 1(3), was performed by the Tukey–Hanning lag-window kernel: $\lambda_m(h) = \frac{1}{2} \cdot \left(1 + \cos\left(\pi \cdot \frac{h}{m} \right) \right)$ for $|h| \leq m$, $\lambda_m(h) = 0$ for $|h| > m$.

The bandwidth b: The estimation of the local Gaussian auto-correlations requires the selection of a bandwidth-vector $\boldsymbol{b} = (b_1, b_2)$, and the plots in this section used $\boldsymbol{b} = (0.6, 0.6)$. A discussion related to the choice of bandwidth can be found in Section S4.

The truncation level m: The value $m = 10$ was used for the truncation level, since it was possible to detect nonlinear dependency structures even for that low a truncation level.

The number of replicates R: The estimated values (means and 90% pointwise confidence intervals) were based on $R = 100$ replicates. Simulations were used for the cases with known parametric models, whereas the bootstrapped-based resampling strategy developed in JT22 was used for the real data example. The relevant details about this resampling strategy are, for the convenience of the reader, included in Section S5.

Numerical convergence: The R-package `localgauss` (see [Berentsen et al. 2014](#)) estimates the local Gaussian cross-correlations $\rho_{k\ell:v}(h)$ and returns them together with an attribute that reveals whether or not the estimation algorithm converged numerically. The m-truncated estimates $\hat{f}_{k\ell:v}^m(\omega)$ inherit the convergence attributes from the estimates $\{\hat{\rho}_{k\ell:v}(h)\}_{h=-m}^m$, and either "NC = OK" or "NC = FAIL" will be added to the plot depending on the convergence status.

Estimation aspects for a given parameter configuration: The estimation of $\hat{f}_{k\ell:v}^m(\omega)$ for a point $v = (v_1, v_2)$ is based on the estimates $\{\hat{\rho}_{k\ell:v}(h)\}_{h=-m}^m$, and it can thus be of interest to see how these estimates depend on the configuration of the abovementioned tuning parameters. A detailed discussion can be found in JT22 [Section 3.2]. The key observation is that the number of pseudo-normalized observations close to the point v will be much lower when v lies in one of the tails. This implies a higher variability of the estimated values $\{\hat{\rho}_{k\ell:v}(h)\}_{h=-m}^m$ for points in the tails, and this is the reason that the estimated pointwise confidence intervals for $\hat{f}_{k\ell:v}^m(\omega)$ are much wider for points in the tails.

Reproducibility and interactive investigations: The R-package `localgaussSpec` contains scripts that can be used to reproduce all the examples and figures encountered in this paper, cf. Section S6.1 in the Supplementary Material for further details.

2.5. Asymptotic Theory for $\hat{f}_{k\ell:v}^m(\omega)$

The asymptotic theory for the local Gaussian cross-spectrum $\hat{f}_{k\ell:v}^m(\omega)$ follows from a few minor adjustments of the asymptotic theory that was developed for the local Gaussian auto-spectra in JT22. A detailed discussion of the theoretical framework (some definitions, and all assumptions and proofs) was relegated to Section S1 in the Supplementary Material. The present discussion will only sketch out the key ideas and then state the asymptotic results related to the local Gaussian cross-spectrum $\hat{f}_{k\ell:v}^m(\omega)$.

The theorems related to the convergence $\hat{f}_{k\ell:v}^m(\omega) \to f_{k\ell:v}(\omega)$ (see Section 2.5.2) require three assumptions. Assumption S1 considers the simultaneous existence of all the local Gaussian cross-correlations $\rho_{k\ell:v}(h)$ at the point v for all lags h, whereas Assumption S2 considers the requirements needed for the individual estimates $\hat{\rho}_{k\ell:v}(h)$ to converge against $\rho_{k\ell:v}(h)$. For the convergence of $\hat{f}_{k\ell:v}^m(\omega)$ against $f_{k\ell:v}(\omega)$, it is necessary to know how m, $\boldsymbol{b}$, and n should interact when the three limits $m \to \infty$, $\boldsymbol{b} \to \boldsymbol{0}^+$, and $n \to \infty$ are taken, and this is specified in Assumption S3.

2.5.1. A Brief Sketch of the Requirements for $\boldsymbol{Y}_t = \left(Y_{1,t}, \ldots, Y_{d,t} \right)$

The local Gaussian spectrum $f_{k\ell:v}(\omega)$ is obtained as the Fourier transform of the local Gaussian cross-correlations $\rho_{k\ell:v}(h)$, and it is thus necessary to add requirements on the k and ℓ components of the multivariate times series $\{\boldsymbol{Y}_t\}_{t\in\mathbb{Z}}$. This will be phrased relative to the bivariate pairs that can be created as different combinations of elements from the univariate marginals $\left\{ Y_{k,t} \right\}_{t\in\mathbb{Z}}$ and $\left\{ Y_{\ell,t} \right\}_{t\in\mathbb{Z}}$. Note that the *folding property* from Item 4 of Algorithm 1 implies that it is sufficient to formulate the assumption based on non-negative values of the lag h.

Definition 4. *For a strictly stationary multivariate time series* $\{\boldsymbol{Y}_t\}_{t\in\mathbb{Z}}$*, with* $\boldsymbol{Y}_t = \left(Y_{1,t}, \ldots, Y_{d,t} \right)$*, and for a selected pair of indices k and ℓ, define the following bivariate pairs from the univariate marginals* $\left\{ Y_{k,t} \right\}_{t\in\mathbb{Z}}$ *and* $\left\{ Y_{\ell,t} \right\}_{t\in\mathbb{Z}}$*:*

$$\boldsymbol{Y}_{k\ell:h:t} := \left[Y_{k,t+h}, Y_{\ell,t} \right]', \qquad h \geq 0, \tag{11a}$$

$$\boldsymbol{Y}_{\ell k:h:t} := \left[Y_{\ell,t+h}, Y_{k,t} \right]', \qquad h \geq 1, \tag{11b}$$

and let $g_{k\ell:h}(\boldsymbol{y}_{k\ell:h})$ *and* $g_{\ell k:h}(\boldsymbol{y}_{\ell k:h})$ *denote the respective probability density functions.*

The basic idea for the construction of $f_{k\ell:v}(\omega)$ is that a point $v = (v_1, v_2)$ should be selected at which for all h the density functions $g_{k\ell:h}(\boldsymbol{y}_{k\ell:h})$ of $\boldsymbol{Y}_{k\ell:h:t}$ will be approximated by $\psi(\boldsymbol{y}_{k\ell:h}; \boldsymbol{\theta}_{k\ell:v:h})$, where ψ is the bivariate Gaussian density function from Equation (4). Note that the parameter vector $\boldsymbol{\theta}$ from Equation (4) here has been replaced with $\boldsymbol{\theta}_{k\ell:v:h}$. The added sub-indices are necessary since the theoretical investigation of $\hat{f}^m_{k\ell:v}(\omega)$ requires a total of $2m + 1$ different parameter vectors to be considered simultaneously.

The local investigation (one for each h) requires a bandwidth vector $\boldsymbol{b} = (b_1, b_2)$ and a kernel function $K(\boldsymbol{w})$, which is used to define $K_{k\ell:h:b}(\boldsymbol{y}_{k\ell:h} - v) := \frac{1}{b_1 b_2} K\left(\frac{y_{k,h} - v_1}{b_1}, \frac{y_{\ell,0} - v_2}{b_2} \right)$, which in turn is used in the function $q_{k\ell:v:h:b}(\boldsymbol{\theta}_{k\ell:v:h})$ defined below:

$$q_{k\ell:v:h:b} := \int_{\mathbb{R}^2} K_{k\ell:h:b}(\boldsymbol{y}_{k\ell:h} - v)\left[\psi(\boldsymbol{y}_{k\ell:h}; \boldsymbol{\theta}_{k\ell:v:h}) - \log \psi(\boldsymbol{y}_{k\ell:h}; \boldsymbol{\theta}_{k\ell:v:h}) g_{k\ell:h}(\boldsymbol{y}_{k\ell:h}) \right] \mathrm{d}\boldsymbol{y}_{k\ell:h}. \tag{12}$$

The target of interest, for a given bandwidth $\boldsymbol{b}$, is to find parameter vectors $\boldsymbol{\theta}_{k\ell:v:h:b}$ that minimize Equation (12). Under suitable assumptions (see Section S1.1 for details), there will exist a unique $\boldsymbol{\theta}_{k\ell:v:h:b}$ when $\boldsymbol{b}$ is small enough. A limiting parameter vector $\boldsymbol{\theta}_{k\ell:v:h:0}$ will then occur when the bandwidth $\boldsymbol{b}$ moves to zero, and the correlation components $\rho_{k\ell:v:h:0}$ from $\boldsymbol{\theta}_{k\ell:v:h:0}$ are then defined to be the local Gaussian cross-correlation $\rho_{k\ell:v}(h)$ between $Y_{k,t+h}$ and $Y_{\ell,t}$. The Fourier transform of all the local Gaussian cross-correlations at the point v, will (when they are absolutely summable, cf. Definition 1) give the local Gaussian cross-spectrum $f_{k\ell:v}(\omega)$.

For a given sample of size n, it is of course not possible to find estimates $\hat{\rho}_{k\ell:v}(h)$ of $\rho_{k\ell:v}(h)$ for all lags h. This implies (cf. Algorithm 1) that it is only possible to create an estimate $\hat{f}^m_{k\ell:v}(\omega)$ of the m-truncated local Gaussian spectrum $f^m_{k\ell:v}(\omega)$.

The convergence results related to an estimate $\hat{\rho}_{k\ell:v}(h)$ of $\rho_{k\ell:v}(h)$ only need to consider how $\boldsymbol{b}$ should depend on n as $\boldsymbol{b} \to \boldsymbol{0}^+$ and $n \to \infty$. The asymptotic theory for the estimates $\hat{f}^m_{k\ell:v}(\omega)$ must also take into account that there is an additional limit $m \to \infty$, and it is thus necessary to specify how m, $\boldsymbol{b}$, and n should interact as $m \to \infty$, $\boldsymbol{b} \to \boldsymbol{0}^+$, and $n \to \infty$. The interested reader can consult Section S1.1 for further details, definitions, and technical assumptions related to this issue. See, in particular, Assumption S3 for the relationship between n, m, and $\boldsymbol{b}$.

2.5.2. Convergence Theorems for $\hat{f}^m_{k\ell:v}(\omega)$, $\hat{\alpha}^m_{k\ell:v}(\omega)$, and $\hat{\phi}^m_{k\ell:v}(\omega)$

Note that Assumptions S1–S3 are given in Section S1.1 in the Supplementary Material. See Section S1.3 for the proofs of the theorems stated below.

Theorem 1. *The estimate $\hat{f}^m_{k\ell:v}(\omega) = \hat{c}^m_{k\ell:v}(\omega) - i \cdot \hat{q}^m_{k\ell:v}(\omega)$ of the local Gaussian cross-spectrum $f_{k\ell:v}(\omega) = c_{k\ell:v}(\omega) - i \cdot q_{k\ell:v}(\omega)$ will, under Assumptions S1–S3, satisfy*

$$\sqrt{n(b_1 b_2)^3/m} \cdot \left(\begin{bmatrix} \hat{c}^m_{k\ell:v}(\omega) \\ \hat{q}^m_{k\ell:v}(\omega) \end{bmatrix} - \begin{bmatrix} c_{k\ell:v}(\omega) \\ q_{k\ell:v}(\omega) \end{bmatrix} \right) \xrightarrow{d} \mathrm{N} \left(\begin{bmatrix} 0 \\ 0 \end{bmatrix}, \begin{bmatrix} \sigma^2_{c|k\ell:v}(\omega) & 0 \\ 0 & \sigma^2_{q|k\ell:v}(\omega) \end{bmatrix} \right), \quad (13)$$

when $\omega \notin \frac{1}{2} \cdot \mathbb{Z} = \left\{ \ldots, -1, -\frac{1}{2}, 0, \frac{1}{2}, 1, \ldots \right\}$, where the variances $\sigma^2_{c|k\ell:v}(\omega)$ and $\sigma^2_{q|k\ell:v}(\omega)$ are given by

$$\sigma^2_{c|k\ell:v}(\omega) = \lim_{m\to\infty} \frac{1}{m} \left(\tilde{\sigma}^2_{k\ell:v}(0) + \sum_{h=1}^{m} \lambda^2_m(h) \cdot \cos^2(2\pi\omega h) \cdot \left\{ \tilde{\sigma}^2_{k\ell:v}(h) + \tilde{\sigma}^2_{\ell k:\tilde{v}}(h) \right\} \right) \quad (14a)$$

$$\sigma^2_{q|k\ell:v}(\omega) = \lim_{m\to\infty} \frac{1}{m} \left(\sum_{h=1}^{m} \lambda^2_m(h) \cdot \sin^2(2\pi\omega h) \cdot \left\{ \tilde{\sigma}^2_{k\ell:v}(h) + \tilde{\sigma}^2_{\ell k:\tilde{v}}(h) \right\} \right), \quad (14b)$$

with $\tilde{\sigma}^2_{k\ell:v}(h)$ and $\tilde{\sigma}^2_{\ell k:\tilde{v}}(h)$ as the asymptotic variances related to the estimates $\hat{\rho}_{k\ell:v}(h)$ and $\hat{\rho}_{\ell k:\tilde{v}}(h)$ (see Theorem S4 in the Supplementary Material for the details).

The local Gaussian quadrature-spectrum is identical to zero when $\omega \in \frac{1}{2} \cdot \mathbb{Z}$, and for those frequencies the following asymptotic result holds under the given assumptions:

$$\sqrt{n(b_1 b_2)^3/m} \cdot \left(\hat{f}^m_{k\ell:v}(\omega) - f_{k\ell:v}(\omega) \right) \xrightarrow{d} \mathrm{N}\left(0, \sigma^2_{c|k\ell:v}(\omega)\right), \qquad \omega \in \frac{1}{2} \cdot \mathbb{Z}. \quad (15)$$

The asymptotic results for the local Gaussian amplitude- and phase-spectra are a direct consequence of Theorem 1 and Brockwell and Davis (1986, Proposition 6.4.3, p. 211).

Theorem 2. *Under Assumptions S1–S3, when $\alpha_{k\ell:v}(\omega) > 0$ and $\omega \notin \frac{1}{2} \cdot \mathbb{Z}$, the estimate $\hat{\alpha}^m_{k\ell:v}(\omega) = \sqrt{\left(\hat{c}^m_{k\ell:v}(\omega)\right)^2 + \left(\hat{q}^m_{k\ell:v}(\omega)\right)^2}$ satisfies*

$$\sqrt{n(b_1 b_2)^3/m} \cdot \left(\hat{\alpha}^m_{k\ell:v}(\omega) - \alpha_{k\ell:v}(\omega) \right) \xrightarrow{d} \mathrm{N}\left(0, \sigma^2_{\alpha}(\omega)\right), \quad (16)$$

where $\sigma^2_{\alpha}(\omega)$ is given relative to $\sigma^2_{c|k\ell:v}(\omega)$ and $\sigma^2_{q|k\ell:v}(\omega)$ (from Equation (14), Theorem 1) as

$$\sigma^2_{\alpha} = \left(c^2_{k\ell:v}(\omega) \cdot \sigma^2_{c|k\ell:v}(\omega) + q^2_{k\ell:v}(\omega) \cdot \sigma^2_{q|k\ell:v}(\omega) \right) / \alpha^2_{k\ell:v}(\omega). \quad (17)$$

Theorem 3. *Under Assumptions S1–S3, when $\alpha_{k\ell:v}(\omega) > 0$ and $\omega \notin \frac{1}{2} \cdot \mathbb{Z}$, the estimate $\hat{\phi}^m_{k\ell:v}(\omega) = \mathrm{args}(\hat{c}^m_{k\ell:v}(\omega) - i \cdot \hat{q}^m_{k\ell:v}(\omega))$ satisfies*

$$\sqrt{n(b_1 b_2)^3/m} \cdot \left(\hat{\phi}^m_{k\ell:v}(\omega) - \phi_{k\ell:v}(\omega) \right) \xrightarrow{d} \mathrm{N}\left(0, \sigma^2_{\phi}(\omega)\right), \quad (18)$$

where $\sigma^2_{\phi}(\omega)$ is given relative to $\sigma^2_{c|k\ell:v}(\omega)$ and $\sigma^2_{q|k\ell:v}(\omega)$ (from Equation (14), Theorem 1) as

$$\sigma^2_{\phi}(\omega) = \left(q^2_{k\ell:v}(\omega) \cdot \sigma^2_{c|k\ell:v}(\omega) + c^2_{k\ell:v}(\omega) \cdot \sigma^2_{q|k\ell:v}(\omega) \right) / \alpha^4_{k\ell:v}(\omega). \quad (19)$$

The asymptotic normality results in Theorems 1–3 do not necessarily help much if computations of pointwise confidence intervals for the estimated local Gaussian estimates are of interest, since in practice it may be unfeasible to find decent estimates of the variances $\sigma^2_{c|k\ell:v}(\omega)$ and $\sigma^2_{q|k\ell:v}(\omega)$ that occur in Theorem 1. The pointwise confidence intervals will,

thus, later in the paper either be estimated based on suitable quantiles obtained by repeated sampling from a known distribution, or they will be based on bootstrapping techniques for those cases where real data have been investigated. Refer to Teräsvirta et al. (2010, Chps. 7.2.5 and 7.2.6) for further details with regard to the need for bootstrapping in such situations.

3. Visualizations and Interpretations

This section will show how different visualizations of the estimates $\hat{f}_{k\ell:v}^{m}(\omega)$ of the m-truncated local Gaussian cross-spectrum $f_{k\ell:v}^{m}(\omega)$ can be used to detect nonlinear dependency structures in a multivariate time series. The plots encountered here are natural extensions of those introduced in JT22, but as $\hat{f}_{k\ell:v}^{m}(\omega)$ is complex-valued, the actual investigation will be based on plots of the corresponding local Gaussian versions of the co-spectrum, quadrature-spectrum, phase-spectrum, and amplitude-spectrum.

A sanity test of the implemented estimation algorithm is presented in Section 3.1, and it is there seen that $\hat{f}_{k\ell:v}^{m}(\omega)$ can detect local periodic structures in an example where a heuristic argument enables the prediction of the anticipated result. Section 3.1 also contains a procedure, based on combined *heatmap and distance* plots, that can help an investigator detect local regions of interest. Section 3.2 applies the local Gaussian machinery to the log-returns of the `EuStockMarkets`-data, and it also contains the results from a GARCH-type model fitted to these log-returns.

A comparison of the results from the original data and the fitted model can reveal to what extent the internal dependency structure of the fitted model actually reflects the dependency structure of the original sample, and this might be of interest with regard to model selection. As in JT22, it will be seen that plots based on $\hat{f}_{k\ell:v}^{m}(\omega)$ can be useful as an exploratory tool, and this approach might detect nonlinear dependencies and periodicities between the variables, which cannot be detected by ordinary cross-spectral analysis.

3.1. Sanity Testing the Implemented Estimation Algorithm

This section will check that the estimates of $f_{k\ell:v}^{m}(\omega)$ behave as expected for a few simple simulated bivariate examples. The approach is similar to the one used in JT22, as the examples are bivariate extensions of those used in that paper, but this paper highlights how an initial investigation of *heatmap and distance* plots can help identify regions that might be of interest to investigate further, cf. Figures 2, 4 and 7. It will be seen that an inspection of the `Co`-, `Quad`-, and `Phase`-plots might be useful exploratory tools for the identification of nonlinear dependency structures in multivariate time series.[4]

The strategy used to create the plots for the simulated data works as follows: First draw a given number of independent replicates from the specified model, and compute $\hat{f}_{k\ell:v}^{m}(\omega)$ and $\hat{f}_{k\ell}^{m}(\omega)$ for each of the replicates. Extract the relevant `Co`-, `Quad`-, and `Phase`-spectra, and use the mean of these m-truncated estimates as estimates of the true (and in general unknown) m-truncated spectra. Suitable upper and lower percentiles of the estimates can be used to produce estimates of the pointwise confidence intervals.

Note that the plots are annotated with the following information: a stamp at the center that specifies the type of spectrum investigated; the numerical convergence status NC in the lower left corner; the truncation level m in the upper left corner; the percentiles of the point v of investigation, and the bandwidth b in the upper right corner; the length n and the number of replicates R in the lower right corner. Later on, for plots based on resampling from a given sample, the plots will also include the block length L in the lower right corner.

3.1.1. Bivariate Gaussian White Noise

For the univariate time series considered in JT22, the sanity testing of the implemented estimation algorithm started with the trivial Gaussian case, since for this case it is known that the local Gaussian auto-spectrum coincides with the ordinary auto-spectrum for Gaussian time series. The local Gaussian cross-spectrum also coincides with the ordinary

cross-spectrum for multivariate Gaussian time series, cf. Lemma 1(1), and it is thus natural to test the implemented algorithms on a bivariate Gaussian example.

The explicit details for this bivariate Gaussian test case are not that relevant for the overall discussion, and those are thus relegated to Section S6.2 in the Supplementary Material. For the present discussion, it suffices to say that the resulting plots are as expected, cf. Figure S6.1. In particular, the estimates $\hat{f}^m_{k\ell:v}(\omega)$ and $\hat{f}^m_{k\ell}(\omega)$ of the $m = 10$ truncated local Gaussian and ordinary cross-spectra (based on 100 replicates) are close to the values they are supposed to have, and the estimated pointwise 90% confidence intervals for $\hat{f}^m_{k\ell:v}(\omega)$ are, as expected, wider when the point v lies in the periphery of the observations.

3.1.2. Bivariate Local Trigonometric Examples

With the exception of the Gaussian case, it is not known how the true local Gaussian cross-spectrum should appear. This implies that is hard to know whether or not the m-truncated `Co`-, `Quad`-, and `Phase`-plots look as they are supposed to, or if what is seen might be the result of an erroneous implementation of the estimation algorithm.

The sanity testing of the implemented estimation algorithm was in JT22 performed by the means of an artificially constructed *local trigonometric* time series, for which it at least could be reasonably argued what the expected outcome should be for some specially designated points v (given a suitable bandwidth b). This approach will here be extended to the bivariate case, that is, *bivariate local trigonometric* time series will be constructed for which it, at some designated points v, can be given a heuristic argument for the expected shape of the estimated local Gaussian `Co`-, `Quad`-, and `Phase`-spectra.

These artificial time series will not satisfy the requirements needed for the asymptotic theory to hold true (as is also the case for standard global spectral analysis), but they can still be used to show how an exploratory tool based on the local Gaussian spectral density can detect local structures that the ordinary spectral density fails to detect. Details related to the two cases investigated in this section are given below, whereas Section S6.5 in the Supplementary Material presents an in-depth explanation of the artificial time series construction and the motivation for the heuristic arguments given in this section.

The reference case: The heuristic argument needed for the bivariate case is identical in structure to the one used in the univariate case, and for the present investigation the reference for the plots later on is based on the following simple bivariate model:

$$Y_{1,t} = \cos(2\pi\alpha t + \phi) + w_{1,t} \text{ and } Y_{2,t} = \cos(2\pi\alpha t + \phi + \theta) + w_{2,t}, \tag{20}$$

where $w_{i,t}$ is Gaussian white noise with mean zero and standard deviation σ, with $w_{1,t}$ and $w_{2,t}$ independent, and where, in addition, it is such that α and θ are fixed for all the replicates, whereas ϕ is drawn uniformly from $[0, 2\pi)$ for each individual replicate. A realization with $\sigma = 0.75$, $\alpha = 0.302$, and $\theta = \pi/3$ is used for the `Co`-, `Quad`-, and `Phase`-plots shown in Figure 1, where 100 independent samples of length 1859 were used to obtain the estimates of the m-truncated spectra and their corresponding 90% pointwise confidence intervals (based on the bandwidth $b = (0.6, 0.6)$). Some useful remarks can be based on this plot, before *the bivariate local trigonometric case* is defined and investigated.

In this particular case, the local Gaussian spectra (in blue)[5] in Figure 1 shares many similarities with the corresponding global spectra (in red).[6] In particular, the peaks of the `Co`- and `Quad`-plots lie, for both the local and global spectra, at the frequency $\omega = \alpha$ (shown in the plots as a vertical line), and the corresponding `Phase`-plots at this frequency lie quite close to the phase adjustment $\theta = \pi/3$ (shown as a horizontal line, positioned with an appropriate sign adjustment). This phenomenon is present both for the point at `10%::10%` and the point at `50%::50%`, i.e., the local information at these quantiles is sufficient to pick up both the spectral peak and the phase adjustment expected from the global model. However, it should be noted that this nice match does not hold for all values of σ. In fact, experiments with different values for σ (plots not included in this paper) indicate that the difference between the local and global spectra becomes larger (in particular for the point

10%::10%) when σ becomes smaller. See JT22 [Appendix G.4.4] for some further comments related to this issue.
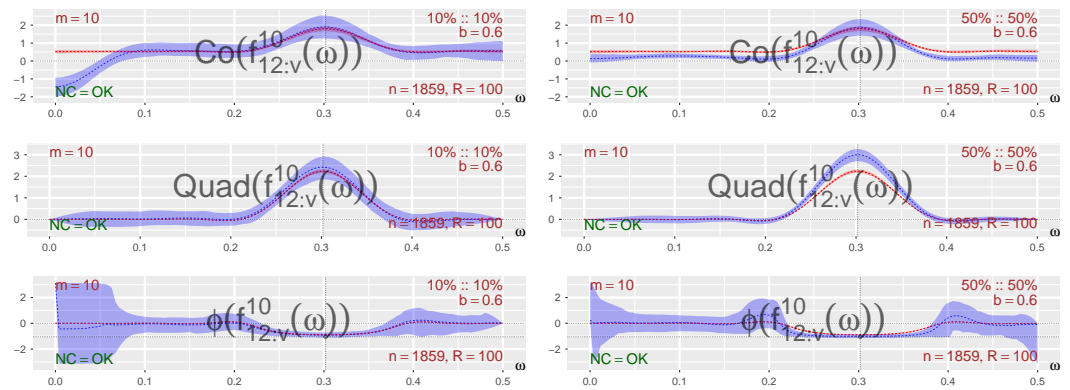

Figure 1. Realization of Equation (20), "bivariate global cosine", with $\sigma = 0.75$, $\alpha = 0.302$, and $\theta = \pi/3$. Similar behavior in lower tail and center, both the global and local Gaussian cross-spectra (truncated at $m = 10$) detect the frequency α and the phase difference θ.

Note that the wide pointwise confidence band observed for ω near 0 in the 10%::10%-Phase-plot in Figure 1 is related to the branch-cut that occurs at $-\pi$ in the definition of the phase-spectrum, cf. Definition 2. The simple algorithm used for the creation of the pointwise confidence intervals was not tweaked to properly cover the case where the majority of the estimates lie in the second and third quadrants of the complex plane, which implies that the Co- and Quad-plots should be consulted instead when the Phase-plot misbehaves in this manner.

The values of the Co- and Quad-plots (for a given frequency ω) are (for each frequency) related to the corresponding values of the Amplitude- and Phase-plots by the following simple relations:

$$-\mathtt{Quad\text{-}plot}/\mathtt{Co\text{-}plot} = \tan(\mathtt{Phase\text{-}plot}), \tag{21a}$$

$$\mathtt{Co\text{-}plot} = \mathtt{Amplitude\text{-}plot} \cdot \cos(\mathtt{Phase\text{-}plot}), \tag{21b}$$

$$-\mathtt{Quad\text{-}plot} = \mathtt{Amplitude\text{-}plot} \cdot \sin(\mathtt{Phase\text{-}plot}), \tag{21c}$$

which follows trivially from the way these spectra are defined relative to Cartesian or polar representations of the complex-valued cross-spectra, cf. Equation (3) and Definition 2. For the example investigated in Figure 1, where the Phase plot is close to $-\pi/3$ at $\alpha = 0.302$, it thus follows that the peak for the Quad-spectrum should be approximately $\sqrt{3}$ times larger than the peak of the Co-spectrum.

The plots encountered later in this paper will be based on the abovementioned Cartesian or polar representations of $\hat{f}_{k\ell:v}^{m}(\omega)$, but it is, in principle, also possible to make plots (and animations) that, for a given value of ω, show the estimated values of $\hat{f}_{k\ell:v}^{m}(\omega)$ as point clouds in the complex plane. A discussion related to this complex-valued plot approach is given in Section S6.4 in the Supplementary Material (see, in particular, Figure S6.2).

The bivariate local trigonometric case: Two bivariate extensions of the artificial *local trigonometric* time series from JT22 [Section 3.3.2] are now considered. The key idea is that an artificial bivariate time series $\left\{ \left(Y_{1,t}, Y_{2,t} \right) \right\}_{t \in \mathbb{Z}}$ can be constructed by the following scheme:

1. Select $r \geq 2$ bivariate time series $\left\{ \left(C_{1,i}(t), C_{2,i}(t) \right) \right\}_{i=1}^{r}$.
2. Select a random variable I with values in the set $\{1, \dots, r\}$, and use this to sample a collection of indices $\{I_t\}_{t \in \mathbb{Z}}$ (that is, for each t, an independent realization of I is taken). Let $p_i := \mathrm{P}(I_i = i)$ denote the probabilities for the different outcomes.

3. Define Y_t by means of the equation

$$Y_{1,t} := \sum_{i=1}^{r} C_{1,i}(t) \cdot \mathbb{1}\{I_t = i\}, \tag{22a}$$

$$Y_{2,t} := \sum_{i=1}^{r} C_{2,i}(t) \cdot \mathbb{1}\{I_t = i\}. \tag{22b}$$

The indicator function $\mathbb{1}\{\cdot\}$ ensures that only one of the bivariate $\left(C_{1,i}(t), C_{2,i}(t)\right)$-components contributes for a given value t, that is, it is also possible to write $\left(Y_{1,t}, Y_{2,t}\right) = \left(C_{1,I_t}(t), C_{2,I_t}(t)\right)$.

The *bivariate local trigonometric* time series (needed for the sanity testing of the implemented estimation algorithm) can now be constructed by selecting r cosine-functions that oscillate around different horizontal baselines L_i, that is,

$$
\begin{aligned}
C_{1,i}(t) &= L_i + A_i(t) \cdot \cos(2\pi\alpha_i t + \phi_i), & i &= 1, \ldots, r, & \text{(23a)} \\
C_{2,i}(t) &= L_i + A_i(t) \cdot \cos(2\pi\alpha_i t + \phi_i + \theta_i), & i &= 1, \ldots, r, & \text{(23b)}
\end{aligned}
$$

where α_i, ϕ_i and θ_i, respectively, represent frequency and phase adjustments occurring in the cosine-function, and where the amplitudes $A_i(t)$ are uniformly distributed in some interval $[a_i, b_i]$. Note that it is assumed that the phases ϕ_i are uniformly drawn (one time for each realization) from the interval between 0 and 2π, whereas the phases θ_i are constants. It is also assumed that the stochastic processes ϕ_i, $A_i(t)$ and I_t are independent of each other.

The auto- and cross-correlations $\rho_{k\ell}(h)$ of $\left\{\left(Y_{1,t}, Y_{2,t}\right)\right\}_{t \in \mathbb{Z}}$, as given by Equations (22) and (23), are functions of L_i and p_i. The auto-correlations $\rho_{11}(h)$ of the marginal time series $\left\{Y_{1,t}\right\}_{t \in \mathbb{Z}}$ were computed in JT22 [Appendix G.4], and the corresponding result for the cross-correlation $\rho_{12}(h)$ is given in Section S6.5 in the Supplementary Material. For the purpose of the present section, it is sufficient to know that it is possible to find parameter configurations for which the global spectrum (based on the pseudo-normalized observations) is rather flat (when truncated at $m = 10$), which implies that it cannot detect the frequencies α_i of the underlying structure. As will be seen in the ensuing simulation experiments, the randomly sampled information for the local versions $C_{1,i}(t)$ and $C_{2,i}(t)$ of Equation (23) may still be sufficient to recover the local periodic and phase structure.

Note that neither $f_{12}(\omega)$ nor $f_{12:v}(\omega)$ are well defined for the *bivariate local trigonometric* times series, but this is not important since it still is possible to predict (cf. Section S6.5 for details) that the m-truncated estimates $\hat{f}_{12:v}^m(\omega)$ for some points v (and a given bandwidth *b*) should resemble Figure 1—and this can be used, cf. Figures 3 and 5, to test the sanity of the implemented estimation algorithm. These plots also reveal that the local Gaussian versions of the `Co`-, `Quad`-, and `Phase`-plots can detect local properties that are undetected by the ordinary version of these spectra.

Parameter setup: The models for the two time series presented in this section are based on an extension of the univariate case seen in JT22, that is, both have $r = 4$ components, the probabilities p_i are given by $(0.05, 1/3 - 0.05, 1/3, 1/3)$, the frequencies α_i are given by $(0.267, 0.091, 0.431, 0.270)$, the baselines L_i are given by the values $(-2, -1, 0, 1)$, and the lower and upper ranges for the uniforms sampling of the amplitudes $A_i(t)$ are, respectively, given by $(0.5, 0.2, 0.2, 0.5)$ and $(1.0, 0.5, 0.3, 0.6)$. Note that L_i and $A_i(t)$ should be selected in order to give a minimal amount of overlap between the different components, cf. Section S6.5 for further details.

The distinction between the two models is due to the selection of the additional phase adjustments θ_i. The model investigated in Figures 2 and 3 has a constant phase adjustment of $\theta = \pi/3$, whereas the model investigated in Figures 4 and 5 has individual phase adjustments, given as $(\theta_1, \theta_2, \theta_3, \theta_4) = (\pi/3, \pi/4, 0, \pi/2)$.

To complete the specification of the setup, note that the 90% pointwise confidence intervals in Figures 3 and 5 all are based on 100 independent samples of length 1859 from the above-described models, and that the bandwidth $b = (0.6, 0.6)$ was used in the computation of the local Gaussian cross-correlations.

Constant phase adjustment: The case where the phase difference $\theta = \pi/3$ was used for all the θ_i is investigated in Figures 2 and 3. This particular example was created in order to test the sanity of the estimation algorithm. The resulting bivariate sample will, by construction, cf. Section S6.5.1 for details, have an ordinary cross-spectrum that looks quite flat (not possible to extract any useful information). Furthermore, cf. the last sentence of the paragraph following the one containing Equation (23), and the heuristic argument given in Section S6.5.2, it is possible to find configurations of bandwidth b and point v for which the resulting local Gaussian spectra should look a bit similar to those encountered in Figure 1 for the three designated points, 10%::10%, 50%::50%, and 90%::90% (which turns out to be the case, cf. Figure 3).

It will, in general, not be known in advance for which points of v that local features might be present, and it is thus necessary to use a strategy that can help identify interesting regions. For this purpose, an adjusted version of the combined *heatmap and distance plots* introduced in JT22 will be used, as seen in Figure 2, where the point v varies along the diagonal. The heatmap part of the plot must be adjusted a bit since the estimated m-truncated local Gaussian cross-spectrum $\hat{f}_{k\ell:v}^m(\omega)$ is a complex-valued entity, and the solution seen in Figure 2 shows a decomposition into the Co- and Quad-plots. The distance parts of these plots are based on the distance function D that is inherited from the complex Hilbert space of Fourier series on the interval $\left[-\frac{1}{2}, \frac{1}{2}\right]$, cf. Section S3.1 for further details. The distance plot can help an investigator see how far the time series of interest is from being i.i.d. observations—and it enables a comparison with the global spectrum that is not possible based only on the heatmap plots.
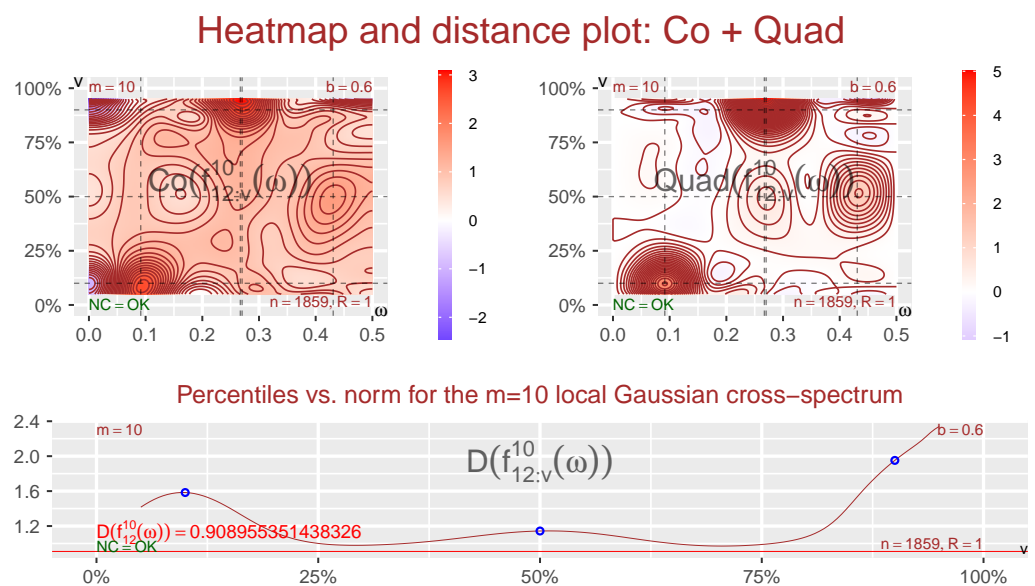
Figure 2. Co- and Quad-heatmaps, distance plot: Sample from the *bivariate local trigonometric* model in Equations (22) and (23), constant phase changes $\theta_i = \pi/3$, $i = 1, 2, 3, 4$. This shows how $\hat{f}_{k\ell:v}^m(\omega)$ varies along the diagonal points v. The frequencies α_i are shown as vertical lines. The points used in Figure 3 are indicated with lines/points.

Note that the heatmap part of the plots can also be based on a polar decomposition into Amplitude- and Phase-plots, but it is easier to digest the information from the Co- and Quad-plots. It is, e.g., easy to see that different local structures are present in the samples investigated by Co- and Quad-based heatmaps in Figures 2 and 4, but this difference is not

that clear when `Amplitude`- and `Phase`-plots are used instead, cf. Section S3.2.2 for further details.

Keep in mind that the pseudo-normalization step in Algorithm 1 implies that the plots related to $\hat{f}^m_{k\ell:v}(\omega)$ only reveal information about the cross-temporal interdependency structures of the sample. An investigator must thus use some supplementing technique in order to extract information from the marginal distributions.

The contours in the heatmap plots seen in Figure 2 reveal that different peaks occur at different combinations of points v and frequencies ω, in agreement with the heuristic argument that motivated the construction of this example. Together with the distance plot part, which shows how much the *m*-truncated local Gaussian cross-spectrum differs from the corresponding *m*-truncated ordinary cross-spectrum, this shows that it could be of interest to take a closer look at the three points investigated in Figure 3.
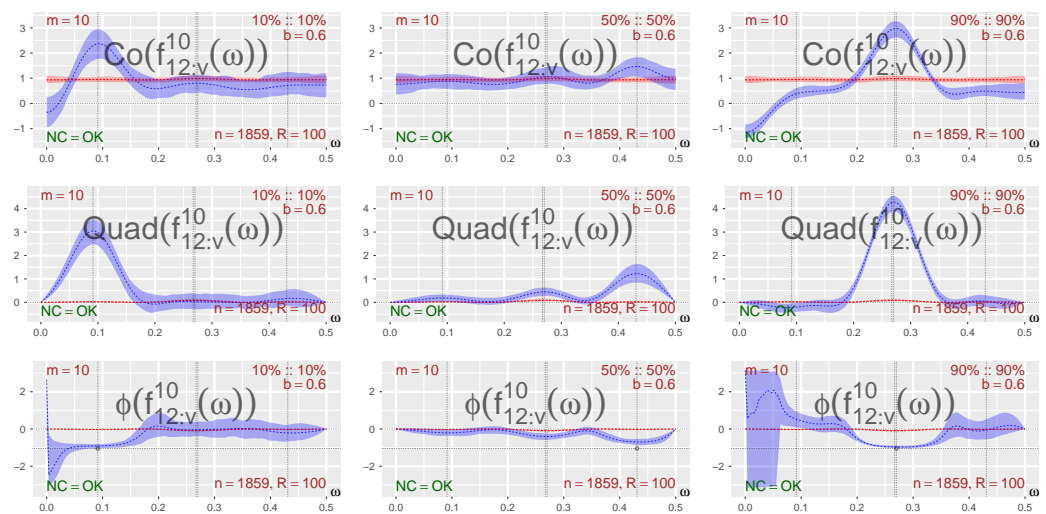
Figure 3. `Co`-, `Quad`-, and `Phase`-plots for three diagonal points: The *bivariate local trigonometric* model in Equations (22) and (23), constant phase changes $\theta_i = \pi/3, i = 1, 2, 3, 4$, shown as sign-adjusted horizontal lines. The frequencies α_i are shown as vertical lines. The local Gaussian spectra detect structures that are not detected by the ordinary spectrum.

The three points investigated in Figure 3 correspond to the function components in Equation (23) with indices $i = 2, 3, 4$. The point that corresponds to the $i = 1$ component would, due to the combination of the low probability p_1 and the placement of the level L_1, lie too far out in the tail to be properly investigated by the present sample size. Note that the $i = 1$ component is included here in order to show that an exploratory approach based on local Gaussian spectra can fail to detect local signals that are much weaker than the dominating ones, cf. the discussion in Section S6.5.

For the points investigated in Figure 3, it seems to be the case that the local Gaussian parts of the `Co`-, `Quad`-, and `Phase`-plots together reveal local properties in accordance with the outcome expected from the knowledge of the generating model—and these local structures are not detected by the ordinary global spectra, which in this case (due to the values used for L_i and p_i) are quite close to being flat (i.e., information about the specified frequencies cannot be extracted from the global spectrum, cf. Section S6.5.1). The left column investigates a point at the lower tail of the diagonal, and it can there be observed that both the `Co`- and `Quad`-plots have a peak close to the leftmost α-value—and the value of the corresponding `Phase`-plot for frequencies close to this α-value lies quite close to the phase difference between the first and second component. A similar situation is present for the three plots shown in the right column, where a point at the upper tail of the diagonal is investigated. Moreover, in accordance with the general discussion related to Equation (21), the peaks of the `Quad`-plots are higher than those of the `Co`-plots in this case due to the phase difference θ that was used in the input parameters.

For the center column of Figure 3, which investigates the point at the center of the diagonal, it can be seen that the Quad- and Phase-plots, in addition to the expected α-value, also detect the presence of the other α-values. The Phase-plot is, for this point, not that close to the expected value, but that situation changes if the truncation is performed at a higher lag than $m = 10$. The center column thus shows the importance of considering a range of values for the truncation level when such plots are investigated. The additional peaks that are detected in the center column are due to *contamination* from the neighboring regions.

Individual phase adjustment: The phase-spectrum contains information about the lead–lag relationship between two time series. In practice, it is conceivable, as indicated in the third paragraph of the Introduction, that the lead–lag relationships are different at extremes; i.e., in the distributional tails, compared to in the center of the series. In fact, such a difference is indicated in the lower row of Figure 9 of the DAX–CAC index pair to be treated in Section 3.2.1 It is therefore of interest to examine the potential of the local Gaussian analysis to pick up such differences for the bivariate local trigonometric example. We therefore regenerated this example with the following individual phase adjustments: $(\theta_1, \theta_2, \theta_3, \theta_4) = (\pi/3, \pi/4, 0, \pi/2)$. The *heatmap and distance* plot seen in Figure 4 reveals, once more, in agreement with the way the example was constructed, that it is natural to look at the three points shown in Figure 5. The horizontal lines seen in the Phase-plots part of Figure 5 show the θ_i-values (adjusted to have the correct sign), and for each of the designated points, the intersection with the relevant vertical α_i-line is highlighted to show the expected outcome based on the knowledge of the model.
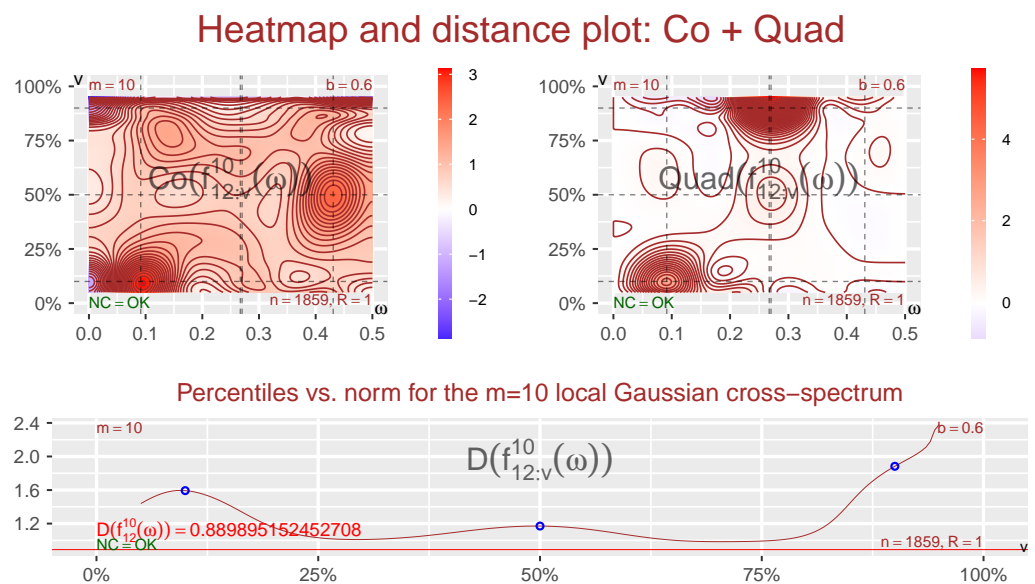
Figure 4. Co- and Quad-heatmaps, distance plot: Sample from the *bivariate local trigonometric* model in Equations (22) and (23), phase changes $(\theta_1, \theta_2, \theta_3, \theta_4) = (\pi/3, \pi/4, 0, \pi/2)$. This shows how $\hat{f}_{k\ell:v}^m(\omega)$ varies along the diagonal points v. The frequencies α_i are shown as vertical lines. The points used in Figure 5 are indicated with lines/points.

The Co-, Quad-, and Phase-plots in Figure 5 behave in accordance with what was observed in Figure 3; in particular, the Phase-plots pick up the individual phase adjustments, giving values that lie close to the expected θ_i-value when the frequency ω is near the corresponding α_i-value. However, there are also changes in the Co- and Quad-spectral plots that can be explained. In fact, the height of the corresponding Co- and Quad-peaks are in accordance with the values of the Phase-plots. In particular, the phase-adjustment is $\theta_2 = \pi/4$ for the point 10%::10%, which implies that the Co- and Quad-peaks should rise approximately to the same height above their respective baselines, which seems to be fairly

close to the observed result. For the points 50%::50% and 90%::90%, the situation is clearer since the respective local frequencies $\theta_3 = 0$ and $\theta_4 = \pi/2$ then imply that only the Co-plot should have a peak for the point 50%::50% and only the Quad-plot should have a peak for the point 90%::90%, again in agreement with the impression based on Figure 5. The global spectrum in Figure 5 (red line) does not reveal any of this local spectral structure.
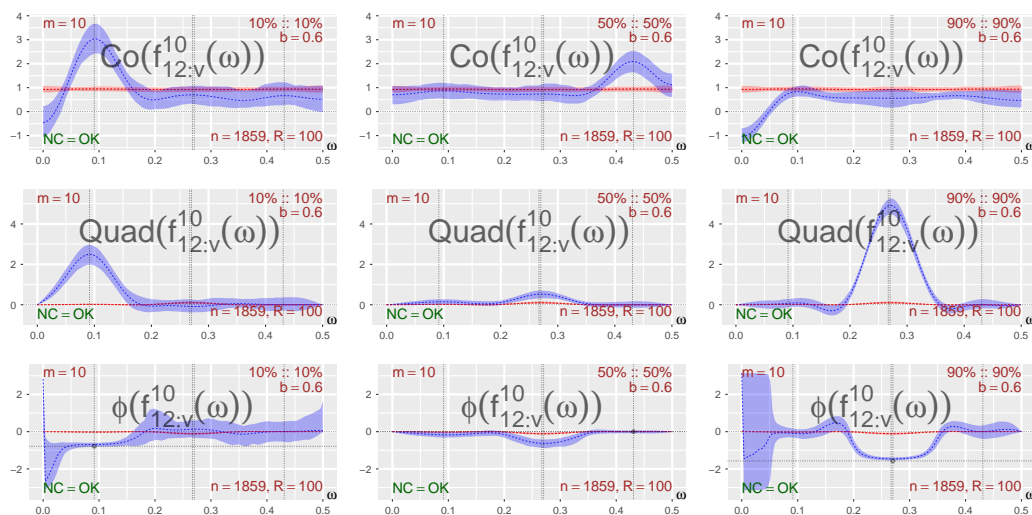

Figure 5. Co-, Quad-, and Phase-plots for three diagonal points: The *bivariate local trigonometric* model in Equations (22) and (23), phase changes $(\theta_1, \theta_2, \theta_3, \theta_4) = (\pi/3, \pi/4, 0, \pi/2)$, shown as sign-adjusted horizontal lines. The frequencies α_i are shown as vertical lines. The local Gaussian spectra detect structures that are not detected by the ordinary spectrum.

The examples investigated in Figures 1–5 do not satisfy the requirements needed for the asymptotic results (both for the global and local cases) to hold true; in particular, the local Gaussian cross-correlations will, in these cases, not be absolutely summable.[7] Despite this, the examples are still of interest since they show that an exploratory tool based on the local Gaussian spectra in this case (with a truncation at $m = 10$) can detect information that is in agreement with what could be expected based on the parameters used in the models for $\left(Y_{1,t}, Y_{2,t}\right)$.

The underlying model will, of course, not be known when a real multivariate time series is encountered, so it is important to estimate the local Gaussian cross-spectrum at a wide range of points v in the plane and a wide range of truncation levels m. This kind of investigation can be performed using the *heatmap and distance* plots seen in Figures 2, 4 and 7.

Even when it might not be obvious how to interpret the results shown in the Co-, Quad-, and Phase-plots, it should be noted that they can be used as an exploratory tool that can detect nonlinear traits in the observations. Moreover, these plots can also be used to investigate if a model fitted to the data contains elements that can mimic the observed features. The recipe for this approach would then be to first select a model, then estimate parameters based on the available sample, and finally use the resulting fitted model to generate independent samples of the same length as the sample. Section 3.2 will show an example of this approach, cf. Figures 9 and 12. Further details related to this are given in Section S5.1.

3.2. A Real Multivariate Time Series and a Poorly Fitted GARCH-Type Model

This section will show how the Co-, Quad-, and Phase-plots can be used as an exploratory tool on a financial dataset, and then it will be seen how this approach also can be used to obtain a visual impression of the quality of a multivariate GARCH-type model fitted to these data. The multivariate time series sample to be considered in this section will be a bivariate subset of the (log-returns of the) tetravariate EuStockMarkets-sample

from the `datasets`-package of R, R Core Team (2020). This dataset was selected since it has a length that should be large enough to justify the assumption that the observed features in the `Co-`, `Quad-`, and `Phase`-plots are not solely there due to small sample variation.

The `EuStockMarkets` contains 1860 daily closing prices collected in the period 1991–1998, from the following four major European stock indices: Germany DAX (Ibis), Switzerland SMI, France CAC, and UK FTSE. The data were sampled in business time, i.e., weekends and holidays were omitted. The log-returns of the `EuStockMarkets` values gives a dataset $(Y_1, Y_2, Y_3, Y_4) = (\text{DAX}, \text{SMI}, \text{CAC}, \text{FTSE})$ that seems natural to model with some multivariate GARCH-type model, and the R-package `rmgarch`, Ghalanos (2022a), was used for that purpose. Note that the present paper only aims to show how this kind of analysis can be performed, so only one very simple model was investigated—which thus gave a rather poor model for the data at hand.

The local Gaussian approach presented here is based on a pseudo-normalization of the marginal distributions. It is thus, in essence, properties of the copula structures of the time series that are revealed here, and a practitioner should supplement this approach by methods that also extract information from the original marginals.

3.2.1. The DAX–CAC Subset of the `EuStockMarkets`-Log-Returns

The log-returns of the bivariate `EuStockMarkets`-subset $(Y_1, Y_3) = (\text{DAX}, \text{CAC})$, of length 1859, will now be investigated.[8] The individual pseudo-normalized traces of these observations are shown in Figure 6, and it will be from these pseudo-normalized observations that the local Gaussian cross-correlations will be computed.
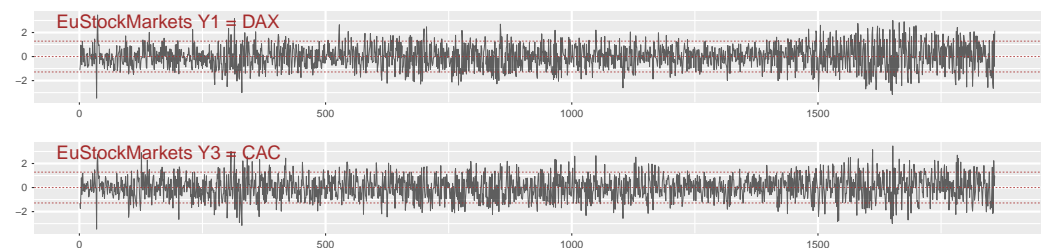
Figure 6. The pseudo-normalized log-reurns of the DAX- and CAC-components from the `EuStockMarkets`-data, which will be investigated further in Figures 7–9.

A local Gaussian investigation requires that some points v must be selected, and thereafter an in-depth investigation can be performed for the selected points. The *heatmap* and *distance* plots can be useful for the task of identifying interesting points along the diagonal, as seen in the Figure 7, where the $m = 10$ truncated `Co-` and `Quad`-spectra are presented for the DAX- and CAC-components of the log-returns of the `EuStockMarkets`-data.

The scale indicators for the heatmap plots in Figure 7 reveal that it is the `Co`-spectrum that dominates, and it is clear that the situation in the tails is different from the one in the center. The asymmetry between the upper and lower tail is not immediately visible from the heatmap part of Figure 7, but it is easy to see from the distance part of the plot that the local Gaussian interdependency structure is stronger in the lower tail.

From the information in Figure 7, it can now be seen that it could be of interest to compare the behavior in the tails with that at the center. The selected points v are highlighted in Figure 7 with the help of added lines (the heatmap part) and circles (the distance part). These points v are the same as those used for Figures 3 and 5, i.e., the diagonal points v whose coordinates correspond to the 10%, 50%, and 90% percentiles of the standard normal distribution.
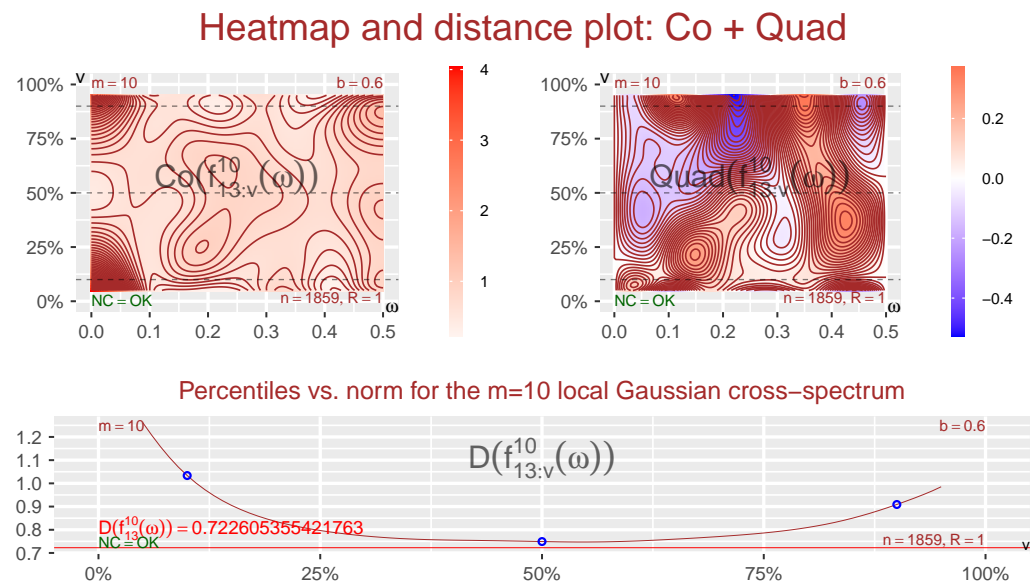
Heatmap and distance plot: Co + Quad

Percentiles vs. norm for the m=10 local Gaussian cross–spectrum

Figure 7. Co- and Quad-heatmaps, distance plot: Based on $\left(Y_{1,t+h}, Y_{3,t}\right)$ with $Y_{1,t+h}$ and $Y_{3,t}$ being, respectively, the DAX- and CAC-components shown in Figure 6. Asymmetric situation, with higher values in the lower tail. The diagonal points used in Figures 8 and 9 are indicated with lines/points.

The estimates of the local Gaussian cross-correlations might degenerate toward $+1$ or -1 if the points v are too far out in the tails. It is thus of interest to check the behavior of these estimates for the given sample before the computationally costly production of pointwise confidence intervals (based on resampling) is undertaken. The result of this investigation of the estimated local Gaussian cross-correlations $\rho_{k\ell:v}(h)$ is seen in Figure 8, where a wide range of values for h is included. The observed values of $\hat{\rho}_{k\ell:v}(h)$ indicate that degeneration of the estimates are not a problem for these points.
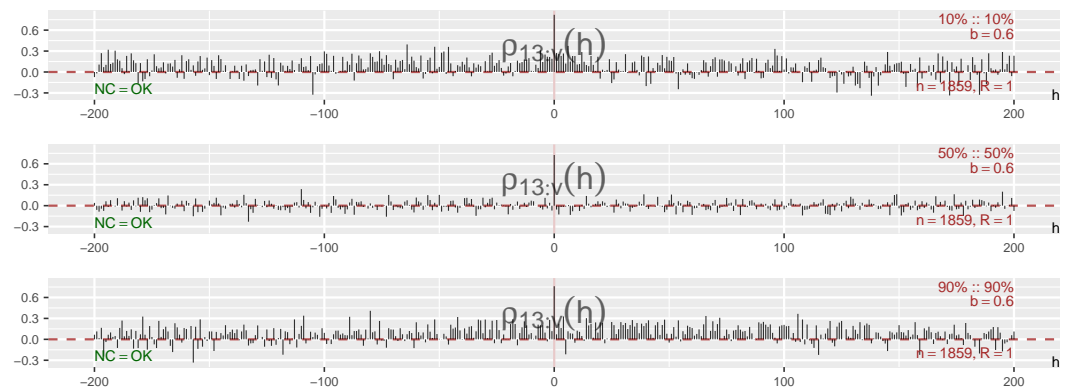
Figure 8. Local Gaussian cross-correlations $\hat{\rho}_{k\ell:v}(h)$ for three diagonal points: between $\left(Y_{1,t+h}, Y_{3,t}\right)$, where $Y_{1,t+h}$ and $Y_{3,t}$ are the DAX- and CAC-components shown Figure 6.

It seems to be the case that the local Gaussian spectrum can detect nonlinear structures even with a rather low truncation level m, but the value $m = 10$ used in this paper is solely selected in order to present a proof-of-principle. For an actual investigation it would be natural to investigate a wider range of different lags h, such as those seen in Figure 8, in order to check what kind of behavior is observed. The interested reader will find a sensitivity analysis of m in Section S3.4, which is based on the values seen in Figure 8.

Figure 9 shows the Co-, Quad-, and Phase-spectra obtained from the $m = 10$ truncated global and local spectra for the three selected diagonal points. A solid red line represents the estimate of the ordinary cross-spectrum $f_{k\ell}^{m}(\omega)$, whereas a solid blue line represents the estimate of the local Gaussian cross-spectrum $f_{k\ell:v}^{m}(\omega)$. The 90% pointwise confidence

intervals were created based on $R = 100$ block-bootstrap replicates using a block length of $L = 25$ for the *circular-index-based block bootstrap for tuples* resampling strategy developed in JT22. Details related to this resampling strategy are given in Section S5 in the Supplementary Material, and the sensitivity analysis in Section S5.3 reveals that the results in this case are stable over a wide range of block lengths.
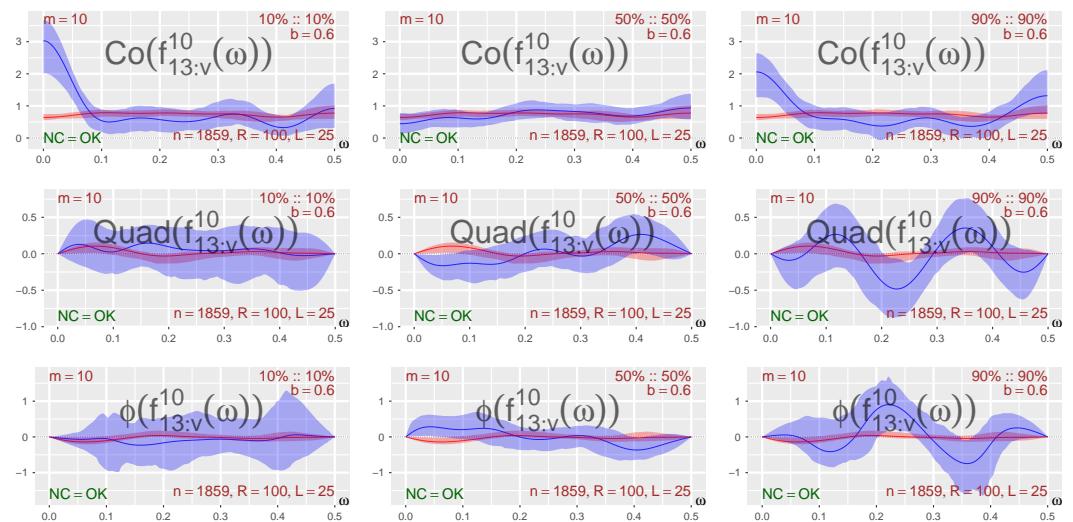
Figure 9. Co-, Quad-, and Phase-plots for three diagonal points: between $\left(Y_{1,t+h}, Y_{3,t}\right)$, where $Y_{1,t+h}$ and $Y_{3,t}$ are the DAX- and CAC-components shown in Figure 6. The local Gaussian spectra detect structures that are not detected by the ordinary spectrum.

The three points considered in Figure 9 all lie on the diagonal, since it seems easier to give an interpretation for those points. In particular, the point 10%::10% represents a situation where the market moves down both in Germany and France, whereas the points 50%::50% and 90%::90% similarly represent cases where the market either is stable or moves up in both countries. For the purpose of the present paper, it suffices to point out that the Co-, Quad-, and Phase-plots of Figure 9 indicate that the data contain nonlinear traits, which in particular is visible in the Co-plot for the point 10%::10% and for all the plots related to the point 90%::90%. It should be noted that the Co-plots at the frequency $\omega = 0$ simply give a weighted sum of the local Gaussian cross-correlations (between $\left(Y_{1,t+h}, Y_{3,t}\right)$) seen in Figure 8, so the Co-plot peaks at $\omega = 0$ for the points 10%::10% and 90%::90% are thus as expected, and the lack of a Co-plot peak at $\omega = 0$ for the point 50%::50% also seems natural in view of Figure 8. It should also be noted that the $\omega = 0$ peak for the 90%::90% Co-plot is lower than the corresponding peak for 10%::10%, but this seems, in this case, to be due to the low truncation level used for the plots, i.e., these two peaks attain approximately the same height when a higher truncation level is applied.

It seems, for the particular parameter configuration that generated the plots in Figure 9, to be the case that the point 90%::90% has the most interesting Quad- and Phase-plots, but again, as noted above, it may be premature to place too much emphasis on this particular plot given the uncertainties involved in the selection of the bandwidth b and the truncation level m. As mentioned before, the effects of changes to the block length L, are quite minimal, cf. the discussion in Section S5.3, in the Supplementary Material.

It should also be noted that a low number of bootstrapped replicates can be a source of small sample variation for the width of the estimated pointwise confidence intervals, and this is important to keep in mind if a minor gap is observed between the pointwise confidence intervals for the local and global spectra. Such gaps could appear or disappear when the algorithm is used to generate new computations based on the same number of bootstrapped replicates, a behavior that, in particular, has been observed for the rightmost peak/trough of the Quad- and Phase-plots at the point 90%::90% in Figure 9. This kind of

ambiguity can be countered by increasing the number of bootstrapped replicates, but that was not carried out for the present example due to the increased computational cost.

The peak at the center of the 90%::90% Phase-plot in Figure 9 seems to be significantly different from the global spectrum, which strengthens the impression that something of interest might be present at that frequency. However, it is important to keep in mind that this impression is based on the present combination of bandwidth and truncation level—and there are, at the moment, no data-driven methods for the selection of these parameters. (A positive phase difference is consistent with the 90%::90% cross-correlation plot in Figure 8, which might indicate that the DAX is leading over CAC when the market is going up. We were not able to find similar evidence in the ordinary spectrum-based literature, but there are other methods for determining lead–lag relationships using comparisons in terms of intensity of jumps of the two series involved, cf. Eyjolfsson and Tjøstheim (2023) for the index pair S&P 500 and Nikkei 225, and Aït-Sahalia et al. (2015) in their investigation of financial contagion.)

Note that the shiny-interface in the R-package localgaussSpec should be used if it is of interest to pursue a further analysis of the local Gaussian spectra of the log-returns of the EuStockMarkets-data, since that enables an interactive investigation that shows how the estimates vary based on different bandwidths b and truncation levels m. Moreover, as discussed in Section S5.3, it is also possible to check how the selection of the block length L for the bootstrap procedure developed in JT22 influences the widths of the pointwise confidence intervals in the Co-, Quad-, and Phase-plots.

3.2.2. A Simple Copula-GARCH-Model Fitted to the EuStockMarkets Log-Returns

It might not be obvious how to interpret the Co-, Quad-, and Phase spectra based on the log-returns of the EuStockMarkets-data, but they do at least provide an approach where nonlinear dependencies might be detected from a visual inspection of the plots.

Furthermore, it is possible to use this as an exploratory tool in order to investigate whether a model fitted to the original data is capable of reproducing nonlinear traits that match those observed for the data. The procedure is straightforward:

1. Fit the selected model to the data.
2. Perform a local Gaussian spectrum investigation based on simulated samples from the fitted model. The parameters should match those used in the investigation of the original data.
3. Compare the plots based on the original data with corresponding plots based on the simulated data from the model. It can be of interest to not only compare the Co-, Quad-, and Phase-plots, but also include plots that show the traces and the estimated local Gaussian auto- and cross-spectra.

For the present case of interest, Item 2 of the list above implies that 100 independent samples of length 1859 will be used as the basis for the construction of the Co-, Quad-, and Phase-plots of the fitted model, and the bandwidth $b = (0.6, 0.6)$ will be used for the estimation of the local Gaussian cross-correlations at the three diagonal points 10%::10%, 50%::50%, and 90%::90%.

The model: The R-package rmgarch was used to fit a simple multivariate GARCH-type model to the log-returns of the EuStockMarkets-data, in order to exemplify the procedure outlined above, i.e., a copula-GARCH-model (cGARCH) with the simplest available univariate models for the marginals[9] was fitted to the data, and the resulting model was then used to produce Figures 10–12. See Section S6.3 for further details. Note that this simple model was selected simply in order to provide a proof-of-principle example for the investigations encountered later on. It was, in particular, of interest, as discussed in Section S5.1, to use a "too-simple model" in order to highlight the ideas related to the local Gaussian sanity testing of parametric models fitted to a sample.

The traces: Figure 10 shows the pseudo-normalized trace of the Y_1- and Y_3-variables for one sample from the tetravariate cGARCH-model, and this can be compared with the corresponding pseudo-normalized trace of the DAX and CAC plot for the pseudo-

normalized log-returns of the `EuStockMarkets`-data (see Figure 6). Obviously, a comparison of one single simulated trace with the trace of the original data might not reveal much, and it should also be noted that it, in general, might be preferable to compare the traces before they are subjected to the pseudo-normalization, since that could detect if the model might fail to produce sufficiently extreme outliers.
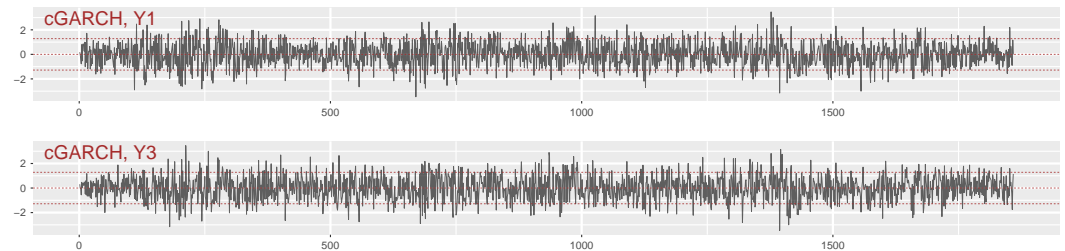

Figure 10. The pseudo-normalized pairs $\left(Y_{1,t+h}, Y_{3,t}\right)$ from a sample from the cGARCH model fitted to the `EuStockMarkets`-data, where $Y_{1,t+h}$ and $Y_{3,t}$ correspond to the DAX- and CAC-components shown in Figure 6.

The local Gaussian correlations: Boxplots, based on the 100 independent estimates of the local Gaussian cross-correlations from the cGARCH-model, are shown in Figure 11. These can be compared with the local Gaussian cross-correlations estimated from the original sample, shown in Figure 8. It should be noted that the computational cost for the production of the boxplots in Figure 11 is substantially larger than the cost for the production of the simpler plots shown in Figure 8, so it is preferable to restrict the attention to a shorter range of lags in Figure 11. Note also that the wide range of lags included in Figure 8 is related to the desire for an initial investigation of the sensitivity of the truncation level m, and it is, as seen here, possible to judge the suitability of the fitted model from the shorter range of lags included in Figure 11.
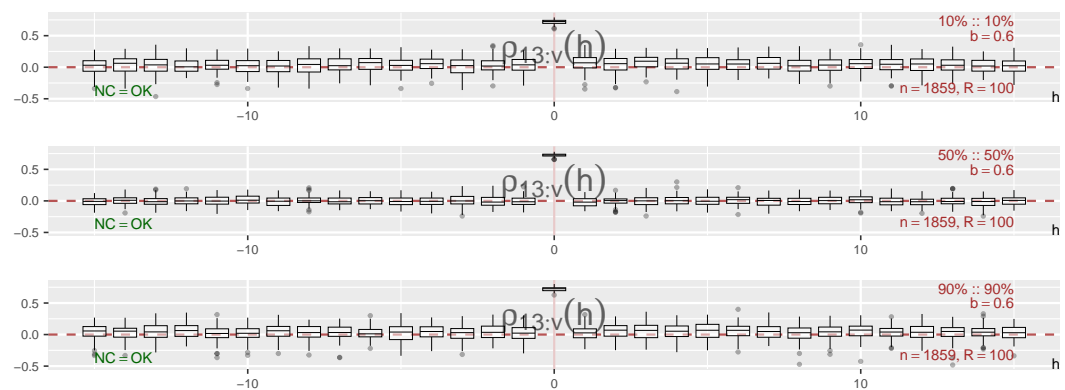
Figure 11. Local Gaussian cross-correlations $\hat{\rho}_{k\ell:v}(h)$ for three diagonal points, based on samples of the pseudo-normalized pairs $\left(Y_{1,t+h}, Y_{3,t}\right)$ from the cGARCH model fitted to the `EuStockMarkets`-data, where $Y_{1,t+h}$ and $Y_{3,t}$ correspond to the DAX- and CAC-components shown in Figure 6. These values were used for the construction of Figure 12.

The impression from the lags included in Figure 11 is that the medians of the estimated local Gaussian cross-correlations for the point `50%::50%` are quite close to zero, whereas the medians for the points `10%::10%` and `90%::90%` are mostly slightly above zero. Almost none of the boxes for the two latter points appear to be positioned in a manner consistent with the desired outcome for a good match with the corresponding estimated values in Figure 8, and it might thus be ample reason to suspect that this cGARCH-model might better be replaced with another model instead.

The `Co`- `Quad`-, and `Phase`-plots: Figure 12 shows the local Gaussian spectra for the same points v and the same configuration of parameters as those used in Figure 9.
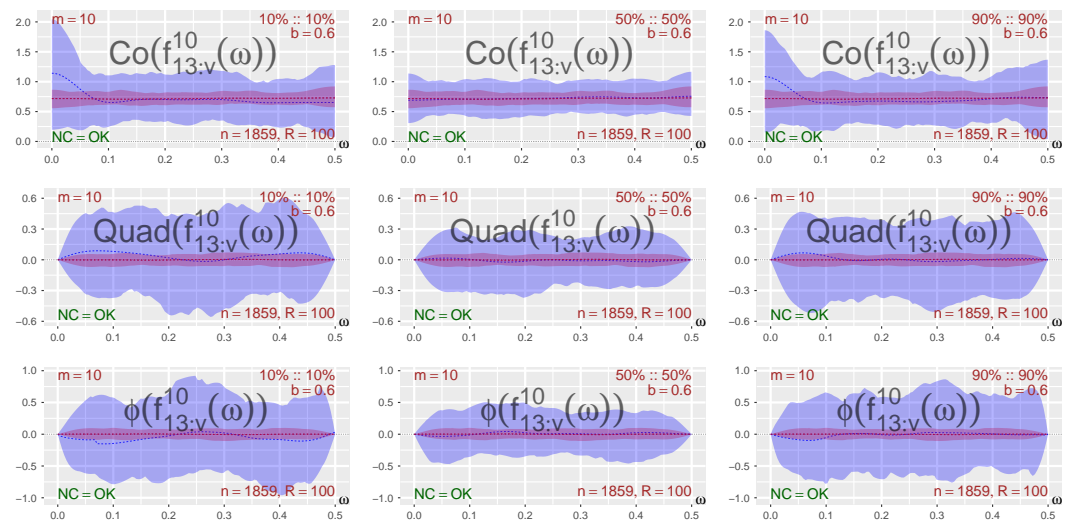
Figure 12. `Co`-, `Quad`-, and `Phase`-plots for three diagonal points, based on samples of the pseudo-normalized pairs $\left(Y_{1,t+h}, Y_{3,t}\right)$ from the cGARCH model fitted to the `EuStockMarkets`-data, where $Y_{1,t+h}$ and $Y_{3,t}$ correspond to the DAX- and CAC-components shown in Figure 6. A comparison with Figure 9 reveals that this model failed to pick up some of the local interdependency structures seen in the `EuStockMarkets`-data.

The `Co`-, `Quad`-, and `Phase`-plots of Figure 12 look as if they could originate from i.i.d. white noise—which comes as no surprise in view of the information about the local Gaussian cross-correlations in Figure 11. The `Co`-plots for the two points `10%::10%` and `90%::90%` show that the estimates of the m-truncated local Gaussian co-spectra, i.e., the blue dashed lines, might have a peak at $\omega = 0$—but the 90% pointwise confidence intervals are too wide to support a claim that these peaks are significant. A further comparison of these two `Co`-plots with the corresponding `Co`-plots in Figure 9 (beware of different scales for the axes) shows that the confidence intervals from Figure 12 are too narrow (at $\omega = 0$) to encompass the peaks observed in Figure 9—which indicates that the selected model might be a rather bad approximation to the log-returns of the `EuStockMarkets`-data. It thus seems advisable to look for some other model instead, a natural conclusion given that no effort whatsoever was made with regard to finding reasonable marginal distributions for the copula-GARCH-model used in this discussion.

It should be noted that a *local Gaussian spectra comparison* of the original data and the fitted model in practice also should include a comparison of the local Gaussian auto-spectra of the marginals, as was performed in JT22. These auto-spectra plots (not included in this paper, but see Section S3.2.1 for some related *heatmap and distance* plots) can provide some additional information useful for the model selection process. In particular, if a model selection algorithm for GARCH-type models is used to pick one marginal model from a given collection of marginal models, then an investigation based on the *local Gaussian auto-spectrum* might reveal if the selected marginal model captures the local traits of the corresponding marginal observations in a reasonable manner.

The preceding discussion was restricted to three diagonal points, but the local Gaussian sanity testing of a fitted parametric model can also be performed for points outside of the diagonal. Such an approach was discussed for the univariate case in JT22 [Appendix F.2], and a natural extension to the multivariate case is presented in Section S5.1 in the online Supplementary Material.

4. Conclusions

The *local Gaussian auto-spectrum* from JT22 can, as seen in this paper, be extended to cover the multivariate case, too—and estimates of the m-truncated *local Gaussian cross-*

spectrum $f^m_{k\ell:v}(\omega)$ can be used to detect nonlinear cross-temporal dependencies between the marginals $Y_{k,t}$ and $Y_{\ell,t}$ of a multivariate time series $\left\{ \mathbf{Y}_t = \left(Y_{1,t}, \cdots, Y_{d,t} \right) \right\}_{t \in \mathbb{Z}}$.

The resulting *local Gaussian approach to spectral analysis* can, in particular, be used to detect if the interdependency structure of the time series under investigation deviates from being Gaussian. For time series whose ordinary auto- and cross-spectra are flat, any peaks and troughs from the *local Gaussian approach* can then be considered indicators of *local nonlinear traits* and *local periodicities*.

The *m*-truncated estimates $\hat{f}^m_{k\ell:v}(\omega)$ can, as discussed for GARCH-models in Section 3.2, also be of interest with regard to *local comparisons* of models fitted to a given sample. In particular, it is possible to use this approach to check if a model fitted to the data can reproduce local traits detected in the original sample, and an investigator can thus use a *local Gaussian sanity test* of the fitted models as a supplement to other model selection methods.

All the examples in this paper can be reproduced with the help of the scripts that are included as a part of the `R`-package `localgaussSpec`. The interested reader can run these scripts, and then use the integrated `shiny`-application from this `R`-package to investigate how the resulting plots change when the tuning parameters in the estimation algorithm are modified. This enables a much deeper inspection of the results than the static plots contained in this paper. These scripts can also be used as templates for new investigations, and they can hopefully help any interested readers to test out the local Gaussian approach to spectral analysis on their own data.

Notes

[1]   This multivariate approach was initiated in the first author's Ph.D. thesis, available at https://bora.uib.no/handle/1956/16950. The version in this paper has been extended with new methods and visualizations that were developed due to review comments related to the univariate theory published in JT22.

[2]   This is due to the way the local Gaussian correlation is defined; see Tjøstheim and Hufthammer (2013) for details.

[3]   The corresponding coordinates are $(-1.28, -1.28)$, $(0, 0)$, and $(1.28, 1.28)$.

[4]   The `Amplitude`-plots are not included here since the interesting details (in most cases) would already have been detected by the other plots.

[5]   If you have a black and white copy of this paper, then read "blue" as "light" and "red" as "dark".

[6]   The dotted lines represent the means of the estimated values, whereas the 90% pointwise confidence intervals are based on the 5% and 95% quantiles of these samples.

[7]   In this respect, the situation is similar to the detection of a pure sinusoidal for the global spectrum.

[8]   The corresponding script in the `R`-package `localgaussSpec` enables an investigation of all the combinations between DAX, SMI, CAC, and FTSE, but only the DAX–CAC subset will be discussed here.

[9]   See Ghalanos (2022a) for details about the cGARCH-model and other options available in the `rmgarch`-package.

[10]   Use `remotes::install_github("LAJordanger/localgaussSpec")` to install the package. See Section S6.1 in the online Supplementary Material for further details.

References

Aït-Sahalia, Yacine, Julio Cacho-Diaz, and Roger J. A. Laeven. 2015. Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics* 117: 585–606. [CrossRef]

Baruník, Jozef, and Tobias Kley. 2019. Quantile coherency: A general measure for dependence between cyclical economic variables. *The Econometrics Journal* 22: 131–52. [CrossRef]

Berentsen, Geir Drage, Tore Selland Kleppe, and Dag Bjarne Tjøstheim. 2014. Introducing `localgauss`, an R Package for Estimating and Visualizing Local Gaussian Correlation. *j-J-STAT-SOFT* 56: 1–18. [CrossRef]

Birr, Stefan, Tobias Kley, and Stanislav Volgushev. 2019. Model assessment for time series dynamics using copula spectral densities: A graphical tool. *Journal of Multivariate Analysis* 172: 122–46. Dependence Models. [CrossRef]

Bollerslev, Tim, and Eric Ghysels. 1996. Periodic Autoregressive Conditional Heteroscedasticity. *Journal of Business & Economic Statistics* 14: 139–51. [CrossRef]

Brillinger, David R. 1965. An Introduction to Polyspectra. *The Annals of Mathematical Statistics* 36: 1351–74. [CrossRef]

Brillinger, David R. 1991. Some history of the study of higher-order moments and spectra. *Statistica Sinica* 1: 24J.

Brillinger, David R., ed. 1984. *The Collected Works of John W. Tukey. Volume I. Time Series: 1949–1964.* Introductory Material by William S. Cleveland and Frederick Mosteller. Wadsworth Statistics/Probability Series. Pacific Grove: Wadsworth.

Brockwell, Peter J., and Richard A. Davis. 1986. *Time Series: Theory and Methods.* New York: Springer, Inc.

Chen, Tianbo, Ying Sun, and Ta-Hsin Li. 2021. A semi-parametric estimation method for the quantile spectrum with an application to earthquake classification using convolutional neural network. *Computational Statistics & Data Analysis* 154: 107069. [CrossRef]

Chung, Jaehun, and Yongmiao Hong. 2007. Model-free evaluation of directional predictability in foreign exchange markets. *Journal of Applied Econometrics* 22: 855–89. [CrossRef]

Ciaburro, Giuseppe, and Gino Iannace. 2021. Machine Learning-Based Algorithms to Knowledge Extraction from Time Series Data: A Review. *Data* 6: 55. [CrossRef]

Eyjolfsson, Heidar, and Dag Tjøstheim. 2023. Multivariate self-exciting jump processes with applications to financial data. *Bernoulli*, *in press*.

Ghalanos, Alexios. 2022a. *rmgarch: Multivariate GARCH Models.* R Package Version 1.3-9. Available online: https://cran.r-project.org/package=rmgarch (accessed on 10 December 2022).

Ghalanos, Alexios. 2022b. *rugarch: Univariate GARCH Models.* R Package Version 1.4-9. Available online: https://cran.r-project.org/package=rugarch (accessed on 10 December 2022).

Hjort, Nils Lied, and M. C. Jones. 1996. Locally parametric nonparametric density estimation. *Annals of Statistics* 24: 1619–47. [CrossRef]

Hong, Yongmiao. 1999. Hypothesis Testing in Time Series via the Empirical Characteristic Function: A Generalized Spectral Density Approach. *Journal of the American Statistical Association* 94: 1201–20. [CrossRef]

Hong, Yongmiao, Jun Tu, and Guofu Zhou. 2007. Asymmetries in Stock Returns: Statistical Tests and Economic Evaluation. *The Review of Financial Studies* 20: 1547–81. [CrossRef]

Jordanger, Lars Arne, and Dag Tjøstheim. 2022. Nonlinear Spectral Analysis: A Local Gaussian Approach. *Journal of the American Statistical Association* 117: 1010–27. [CrossRef]

Klimko, Lawrence A., and Paul I. Nelson. 1978. On Conditional Least Squares Estimation for Stochastic Processes. *Annals of Statistics* 6: 629–42. [CrossRef]

Künsch, Hans R. 1989. The Jackknife and the Bootstrap for General Stationary Observations. *The Annals of Statistics* 17: 1217–41. [CrossRef]

Li, Ta-Hsin. 2020. From zero crossings to quantile-frequency analysis of time series with an application to nondestructive evaluation. *Applied Stochastic Models in Business and Industry* 36: 1111–30. [CrossRef]

Li, Ta-Hsin. 2021. Quantile-frequency analysis and spectral measures for diagnostic checks of time series with nonlinear dynamics. *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 2: 270–90. [CrossRef]

Li, Ta-Hsin. 2022a. *Quantile Fourier Transform, Quantile Series, and Nonparametric Estimation of Quantile Spectra*. *Preprint*. [CrossRef]

Li, Ta-Hsin. 2022b. Quantile-Frequency Analysis and Deep Learning for Signal Classification. *Preprint*. [CrossRef]

Li, Zeda. 2023. Robust conditional spectral analysis of replicated time series. *Statistics and Its Interface* 16: 81–96. [CrossRef]

Otneim, Håkon, and Dag Tjøstheim. 2017. The locally Gaussian density estimator for multivariate data. *Statistics and Computing* 27: 1595–616. [CrossRef]

Politis, Dimitris N., and Joseph P. Romano. 1992. A General Resampling Scheme for Triangular Arrays of α-Mixing Random Variables with Application to the Problem of Spectral Density Estimation. *The Annals of Statistics* 20: 1985–2007. [CrossRef]

R Core Team. 2020. *R: A Language and Environment for Statistical Computing*. Vienna: R Foundation for Statistical Computing.

Teräsvirta, Timo, Dag Tjøstheim, and Clive William John Granger. 2010. *Modelling Nonlinear Economic Time Series*. Oxford: Oxford University Press.

Tjøstheim, Dag, and Karl Ove Hufthammer. 2013. Local Gaussian correlation: A new measure of dependence. *Journal of Econometrics* 172: 33–48. [CrossRef]

Tjøstheim, Dag, Håkon Otneim, and Bård Støve. 2021. *Statistical Modeling Using Local Gaussian Approximation*. Cambridge: Academic Press.

Tukey, John W. 1959. An introduction to the measurement of spectra. In *Probability and Statistics, The Harald Cramér Volume*. Edited by Ulf Grenander. Stockholm: Almqvist and Wiksell, pp. 300–30.

Zhao, Zifeng, Peng Shi, and Zhengjun Zhang. 2022. Modeling multivariate time series with copula-linked univariate d-vines. *Journal of Business & Economic Statistics* 40: 690–704. [CrossRef]