

Higher-order accurate, positive semi-definite estimation of large-sample covariance and spectral density matrices

Dimitris N. Politis*

Dept. of Mathematics and Economics

University of California, San Diego

La Jolla, CA 92093-0112, USA

dpolit@ucsd.edu

March 2009

(this is a revised version of Discussion Paper 2005-03)

Abstract

A new class of large-sample covariance and spectral density matrix estimators is proposed based on the notion of flat-top kernels. The new estimators are shown to be higher-order accurate when higher-order accuracy is possible. A discussion on kernel choice is presented as well as a supporting finite-sample simulation. The problem of spectral estimation under a potential lack of finite fourth moments is also addressed. The higher-order accuracy of flat-top kernel estimators typically comes at the sacrifice of the positive semi-definite property. Nevertheless, we show how a flat-top estimator can be modified to become positive semi-definite (even strictly positive definite) while maintaining its higher-order accuracy. In addition, an easy (and consistent) procedure for optimal bandwidth choice is given; this procedure estimates the optimal bandwidth associated with each individual element of the target matrix, automatically sensing (and adapting to) the underlying correlation structure.

*Research partially supported by NSF grants SES-04-18136 and DMS-07-06732. Many thanks are due to Arthur Berg and Tucker McElroy for numerous helpful interventions, to Peter Robinson for a critical reading and suggestions of some key early references, and to Dimitrios Gatzouras for his help with the proof of Lemma 10.1. The S+ software for the practical computation of the different spectral density estimators was compiled with the invaluable help of Isheeta Nargis and Arif Dowla of www.stochasticlogic.com, and is now publicly available from: www.math.ucsd.edu/~dpolit/SOFT/SfunctionsFLAT-TOPS.html.

HIGHER-ORDER ACCURATE, POSITIVE SEMI-DEFINITE ESTIMATION OF LARGE-SAMPLE COVARIANCE AND SPECTRAL DENSITY MATRICES

1 Introduction

Many applications of multivariate time series analysis involve the nonparametric estimation of spectral density matrices. For example, the large-sample covariance matrix of the sample mean of a stationary sequence equals 2π times its spectral density matrix evaluated at the origin. Pioneering work in multivariate spectral estimation was conducted by E.J. Hannan, E. Parzen, M. Rosenblatt, D. Brillinger, and other prominent statistical researchers in the 1950s and 60s. See, e.g., the papers by Hannan (1957, 1958), Parzen (1957, 1961), Priestley (1962), Brillinger and Rosenblatt (1967), as well as the book-length treatments in Hannan (1970), Brillinger (1981), Priestley (1981), and Rosenblatt (1985) that contain a number of additional references.

The subject was revived more recently in the time series econometrics literature where typical applications—such as hypothesis tests from generalized method of moments estimation (Hansen (1982)) or general dynamic models (Gallant and White (1988))—require accurate estimation of large-sample covariance matrices that is robust to autocorrelation and heteroskedasticity. A general theory towards heteroskedasticity and autocorrelation consistent (HAC) covariance matrix estimation was put forth in influential papers by Newey and West (1987) and Andrews (1991); see also related work of Gallant (1987), Andrews and Monahan (1992), Hansen (1992), and Newey and West (1994).

Nevertheless, the current state-of-the-art seems to be lacking in three respects:

- (a) The accuracy of the HAC covariance estimators is often suboptimal as their rate of convergence is $T^{2/5}$ even in situations when higher-order accuracy is possible, e.g., a rate closer to $T^{1/2}$; see Samarov (1977).
- (b) The problem of optimal bandwidth choice for the HAC estimators has not been conclusively addressed. For example, the ‘plug-in’ procedure of Andrews (1991) will not give consistent estimation of the optimal bandwidth unless the parametric model used to estimate the ‘plug-in’ values holds true. On the other hand, cross-validation methods may give consistent bandwidth estimates but their consistency is typically achieved at a very slow rate; see e.g. Robinson (1991) and the references therein.
- (c) The existing literature focuses on obtaining a single optimal bandwidth, common for

estimating all elements of the target matrix; this is suboptimal as each element of the target matrix generally comes with its own individual optimal bandwidth.

In this paper we address the above three issues. A new class of HAC covariance matrix and spectral density matrix estimators is proposed based on the notion of a *flat-top* kernel defined in Politis (2001) that is a generalization of the trapezoidal kernels of Politis and Romano (1995). The new estimators are shown to be higher-order accurate when higher-order accuracy is possible; a discussion on kernel choice is presented as well as a supporting finite-sample simulation.

The higher-order accuracy of flat-top kernel estimators typically comes at the sacrifice of the positive semi-definite property. Nevertheless, we show how a flat-top estimator can be modified to become positive semi-definite (even strictly positive definite) while maintaining its higher-order accuracy. In addition, it is shown that there is an easy (and consistent) procedure for optimal bandwidth choice for flat-top kernel HAC estimators; this procedure estimates the optimal bandwidth associated with each individual element of the target matrix, automatically sensing (and adapting to) the underlying correlation structure.

Since estimation of the large-sample covariance matrix of a sample mean or generalized method of moments estimator is tantamount to estimation of a spectral density matrix evaluated at the origin, the paper treats the more general framework of higher-order accurate, positive semi-definite estimation of spectral density matrices.

The structure of the paper is as follows. In the next section, the flat-top estimators of a spectral density matrix are defined, and a general theorem on their asymptotic accuracy is given. Section 3 addresses the difficult problem of spectral estimation under a potential lack of finite fourth moments; surprisingly, it is shown that the flat-top estimators retain—for the most part—their higher-order accuracy. Section 4 introduces a modification of the flat-top matrix estimators that results into an estimator that is positive semi-definite (even positive definite—if so desired) while retaining the estimators' higher-order accuracy.

Section 5 discusses some interesting kernels of the flat-top family, while Section 6 is devoted to the issue of data-dependent bandwidth choice. An empirical rule for choosing the bandwidth of a flat-top kernel is given extending the trapezoidal kernel bandwidth choice of Politis (2003); a general asymptotic theorem shows the bandwidth choice rule works by automatically adapting to the underlying (unknown) correlation structure *even* in the lack of finite fourth moments. Section 7 presents some finite-sample simulations complementing our asymptotic results where the high accuracy and rate of convergence of the flat-top

estimators are manifested in practice; some simulations on the problem of hypothesis testing are also reported. Finally, Appendix A addresses in detail the set-up of large-sample HAC covariance matrix estimation that is of interest in econometric applications; Lemma 8.1, in particular, is important as it allows the large-sample covariance matrix estimation to be cast in the framework of spectral density matrix estimation. Appendix B discusses the interesting robustness of plug-in bandwidth estimators in the presence of infinite fourth moments. All technical proofs are placed in the Technical Appendix.

2 Spectral density matrix estimation

Consider observations V_1, \dots, V_T from a second-order stationary d -variate time series $\{V_t, t \in \mathbb{Z}\}$ possessing mean zero¹ and autocovariance matrix sequence $\Gamma(j)$ defined as

$$\Gamma(j) = EV_t V_{t+j}' \quad \text{for } j \geq 0, \quad \text{and } \Gamma(j) = \Gamma(-j)' \quad \text{for } j < 0. \quad (1)$$

Under typical weak dependence conditions—see e.g. Hannan (1970), Brillinger (1981), Brockwell and Davis (1991), or Hamilton (1994)—the spectral density matrix evaluated at point w is defined as

$$F(w) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \Gamma(k) e^{-ikw} \quad (2)$$

where $i = \sqrt{-1}$. The $d \times d$ matrix $F(w)$ is positive semi-definite and Hermitian for any $w \in [-\pi, \pi]$ but note that its off-diagonal elements are, in general, complex-valued; $F_{jk}(w)$ will denote the (j, k) element of $F(w)$. Nevertheless, $F(0)$ has all its elements real-valued, and it is easy to see that $F(0) = \Omega/(2\pi)$ where Ω is defined as

$$\Omega = \lim_{T \rightarrow \infty} \frac{1}{T} \sum_{k=1}^T \sum_{j=1}^T EV_k V_j'. \quad (3)$$

Hence, accurate estimation of $F(0)$ is tantamount to accurate estimation of Ω which represents the large-sample covariance matrix of the sample mean; see Remark 3.1 for details. In what follows, we will consider the more general problem of estimation of $F(w)$ at an arbitrary (fixed) point $w \in [-\pi, \pi]$; since w will be fixed, the short-hand notation F will be used to denote $F(w)$, and F_{jk} will denote the (j, k) element of F .

¹The case of unknown mean is discussed in Remark 3.1, and then again in Section 7 and Appendix A.

To describe our new spectral matrix estimator, we need the notion of a ‘flat-top’ kernel. The general family of flat-top kernels was introduced in Politis (2001). Its typical member is $\lambda_{g,c}(x)$ where

$$\lambda_{g,c}(x) = \begin{cases} 1 & \text{if } |x| \leq c \\ g(x) & \text{else;} \end{cases} \quad (4)$$

here $c > 0$ is a parameter, and $g : \mathbb{R} \rightarrow [-1, 1]$ is a symmetric function, continuous at all but a finite number of points, and satisfying $g(c) = 1$, and $\int_{\mathbb{R}} g^2(x) dx < \infty$. The kernel $\lambda_{g,c}(x)$ is ‘flat’, i.e., constant, over the region $[-c, c]$, hence the name flat-top.

If g is such that $g(x) = 0$ for $|x| \geq$ some x_0 , then the kernel $\lambda_{g,c}(x)$ has a hard cut-off. The simplest representative of such a flat-top kernel has a trapezoidal shape defined as

$$\lambda_{TR,c}(x) = \begin{cases} 1 & \text{if } |x| \leq c \\ \frac{|x|-1}{c-1} & \text{if } c < |x| \leq 1 \\ 0 & \text{else} \end{cases} \quad (5)$$

with $c \in (0, 1]$, i.e., the function g performs a linear interpolation between the values $g(c) = 1$ and $g(1) = 0$. The trapezoidal kernel’s favorable properties were documented in Politis and Romano (1995). The trapezoid may be seen as a cross between the ‘truncated’ kernel defined as $\kappa_{trunc}(x) = 1$ if $|x| \leq 1$ and $\kappa_{trunc}(x) = 0$ else, and the well-known triangular Bartlett kernel $\kappa_B(x) = (1 - |x|)^+$. As a matter of fact, $\lambda_{TR,c}(x)$ tends to $\kappa_{trunc}(x)$ and/or $\kappa_B(x)$ by letting c tend to 1 or 0 respectively. Here, the notation $(y)^+$ indicates the positive part of y , i.e., $(y)^+ = \max(y, 0)$.

Let S be a $d \times d$ matrix of bandwidth parameters with (j, k) element denoted by S_{jk} . As usual, S is thought of as a function of T although this dependence will not be explicitly denoted. The estimator of F that we will consider is \hat{F} with (j, k) element given by:

$$\hat{F}_{jk} = \frac{1}{2\pi} \sum_{m=-T}^T \lambda_{g,c}(m/S_{jk}) \hat{\Gamma}_{jk}(m) e^{-imw} \quad (6)$$

where $\lambda_{g,c}$ is some chosen member of the flat-top family, and $\hat{\Gamma}_{jk}(m)$ is the (j, k) element of the sample autocovariance matrix $\hat{\Gamma}(m)$ defined as

$$\hat{\Gamma}(j) = \frac{1}{T} \sum_{t=1}^{T-j} V_t V'_{t+j} \quad \text{for } j \geq 0; \quad \hat{\Gamma}(j) = \hat{\Gamma}(-j)' \quad \text{for } j < 0, \quad (7)$$

and $\hat{\Gamma}(j) = 0$ for $|j| \geq T$. Note that the dependence of \hat{F}_{jk} on the chosen $\lambda_{g,c}$ is not explicitly denoted.

The favorable large-sample properties of \hat{F} are manifested in the following theorem.

Theorem 2.1 *Assume conditions strong enough to ensure that²*

$$\text{Var}(\hat{F}_{jk}) = O(S_{jk}/T) \quad \text{for any fixed } j, k; \quad (8)$$

Then, for each combination of j and k , the following are true.

(i) *If $\sum_{m=-\infty}^{\infty} |m|^r |\Gamma_{jk}(m)| < \infty$ for some real number $r \geq 1$, then letting S_{jk} proportional to $T^{1/(2r+1)}$ yields*

$$\hat{F}_{jk} = F_{jk} + O_P(T^{-r/(2r+1)}).$$

(ii) *If $|\Gamma_{jk}(m)| \leq Ce^{-am}$ for some constants $C, a > 0$, then letting $S_{jk} \sim A \log T$, for some appropriate constant A , yields*

$$\hat{F}_{jk} = F_{jk} + O_P\left(\frac{\sqrt{\log T}}{\sqrt{T}}\right);$$

as usual, the notation $A \sim B$ means $A/B \rightarrow 1$.

(iii) *If $\Gamma_{jk}(m) = 0$ for $|m| > \text{some } q$, then letting $S_{jk} = \max(\lceil q/c \rceil, 1)$, yields³*

$$\hat{F}_{jk} = F_{jk} + O_P\left(\frac{1}{\sqrt{T}}\right);$$

here $\lceil x \rceil$ is the ‘ceiling’ function, i.e., the smallest integer larger or equal to x .

The conditions of the three parts of Theorem 2.1 are usual conditions of weak dependence. For example, if $\Gamma_{jj}(m) = 0$ for $|m| > \text{some } q$, then the j th coordinate of V_t , say $V_t^{(j)}$, can be thought to follow a Moving Average (MA) model of order q . Similarly, the condition $|\Gamma_{jj}(m)| \leq Ce^{-am}$ is satisfied if $V_t^{(j)}$ follows a stationary ARMA (p, q) model, i.e., AutoRegressive with Moving Average residuals; see e.g. Brockwell and Davis (1991). The polynomial decay in condition (i) is a worst-case scenario; suffices to note that in order to even define the spectral density of $V_t^{(j)}$ the typical condition is $\sum_{m=-\infty}^{\infty} |\Gamma_{jj}(m)| < \infty$, i.e., $r = 0$ in condition (i).

Theorem 2.1 demonstrates the improvement in rate of convergence afforded by the use of flat-top kernels as compared to the $O_P(T^{-2/5})$ error associated with traditional second-order kernels. Most importantly, flat-top kernels are seen to *attain* the lower bounds for the

²There exist different sets of conditions sufficient for eq. (8). Assumption A of Andrews (1991) is such a condition based on summability of fourth cumulants; different conditions based on moment and mixing assumptions are also available, see e.g. Hannan (1970), Brillinger (1981), or Brockwell and Davis (1991).

³Taking the maximum of $\lceil q/c \rceil$ and 1 is done to cover the possibility that $q = 0$.

order of magnitude of the error of a quadratic spectral density estimator under the three cases of Theorem 2.1; these lower bounds are due to Samarov (1977).

Note that Theorem 2.1 not only gives the rate of convergence of \hat{F}_{jk} to F_{jk} , but at the same time it suggests the optimal values of the bandwidth parameter S_{jk} ; here optimality is meant with respect to optimizing the rate of convergence of \hat{F}_{jk} . As is apparent, the optimal S_{jk} crucially depends on the rate of decay of $\Gamma_{jk}(m)$ as m increases. If we had some reason to believe that the rate of decay of $\Gamma_{jk}(m)$ is the *same* for all j, k , then we could let S_{jk} equal some common value s_T , in which case our estimator would take the familiar simple form

$$\hat{F}_{simple} = \frac{1}{2\pi} \sum_{m=-T}^T \lambda_{g,c}(m/s_T) \hat{\Gamma}(m) e^{-imw}; \quad (9)$$

letting $w = 0$, it is seen that the above is of the same exact form as the Newey-West (1987) and Andrews (1991) estimator $\hat{\Omega}$ given in eq. (32). Nevertheless, there is typically no reason to believe that the rate of decay of $\Gamma_{jk}(m)$ is common for all j, k . Thus, \hat{F} is generally preferable to \hat{F}_{simple} .

To elaborate, consider the following example. Let $V_t = (V_t^{(1)}, V_t^{(2)}, V_t^{(3)})'$ where $V_t^{(1)}$ follows an MA(q_1) model, $V_t^{(2)}$ follows an MA(q_2) model independent of $V_t^{(1)}$, and $V_t^{(3)} = V_{t-L}^{(2)}$ for all t . Suppose that the trapezoidal kernel $\lambda_{TR,1/2}(x)$ is used, i.e., $c = 1/2$. Then, Theorem 2.1 (iii) suggests the following optimal bandwidth parameters: $S_{11} = 2q_1$, $S_{22} = 2q_2$, $S_{33} = 2q_2$, $S_{12} = S_{21} = 1$, $S_{13} = S_{31} = 1$, and $S_{23} = S_{32} = 2(q_2 + L)$.

Parts (ii), (iii)—as well as part (i) with $r > 2$ —of Theorem 2.1 show that the rate of convergence of \hat{F} is superior to the Newey-West (1987) estimator based on Bartlett’s kernel, as well as to all second order kernel estimators considered by Andrews (1991); the Newey-West (1987) estimator only achieves a rate of convergence of $T^{1/3}$, while the second order kernels (including the optimal quadratic spectral window) achieve a rate of convergence of $T^{2/5}$.

Remark 2.1 If a chosen bandwidth happens not to be small as compared to the sample size, then the standard asymptotics—such as eq. (8)—might not provide accurate approximations, and the so-called “fixed- b ” asymptotics of Kiefer and Vogelsang (2002, 2005), and higher-order expansions of Hashimzade and Vogelsang (2008) are a valuable alternative. There is no inherent discrepancy between the notion of flat-top kernels and “fixed- b ”

asymptotics.⁴ Indeed, the latter could—in principle—be used in connection with flat-top kernels but it seems that the improvements will be marginal (if at all). The reason is that for small bandwidths, the “fixed- b ” asymptotics coincide with the traditional approximations, and that flat-top kernels are characterized by small (even ultra-small) optimal bandwidths; see e.g. the logarithmic and constant optimal bandwidths in Theorem 2.1 (ii) and (iii).

Remark 2.2 The asymptotic normality of $\sqrt{T/S_{jj}}(\hat{F}_{jj} - E\hat{F}_{jj})$ has been shown under a variety of weak dependence and moment conditions; see e.g. Hannan (1970), Brillinger (1981), Brockwell and Davis (1991), Rosenblatt (1984), Politis and Romano (1995), or Shao and Wu (2007). The limiting variance of $\sqrt{T/S_{jj}}\hat{F}_{jj}$ is then given by the L_2 norm of the employed kernel (doubled when w is an integer multiple of 2π). Similarly, $\sqrt{T/S_{jk}}(\hat{F}_{jk} - E\hat{F}_{jk})$ can be shown to have a limiting *complex* normal distribution; see the first three references listed above.

3 Spectral estimation in the absence of finite fourth moments

As mentioned in the last section, eq. (8) is typically satisfied for kernel estimators such as \hat{F} . Nevertheless, it has been conjectured that some financial time series might not possess finite fourth moments; see Davis and Mikosch (2000), Hall and Yao (2003) or Politis (2004) for a discussion. But if the series $\{V_t\}$ does not possess finite fourth moments, then $Var(\hat{F}_{jk})$ is not well-defined. For this reason, it is convenient to also define the correlation/cross-correlation matrix $\rho(m)$ with (j, k) element given by $\rho_{jk}(m) = \Gamma_{jk}(m)/\sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}$, and estimated by $\hat{\rho}_{jk}(m) = \hat{\Gamma}_{jk}(m)/\sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)}$. We can then define the normalized spectral density matrix evaluated at point w as

$$f(w) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \rho(k)e^{-ikw}; \quad (10)$$

the short-hand notation f will again be used to denote $f(w)$, and f_{jk} will denote the (j, k) element of f . The corresponding flat-top kernel estimator of f is \hat{f} with (j, k) element given

⁴The “steep-origin” kernels of Phillips, Sun, and Jin (2006) are competitors to the “fixed- b ” asymptotics but the underlying idea is the same, i.e., better approximations when the bandwidth happens to be large; here the kernel is raised to a power instead of being re-scaled by the bandwidth parameter. Note though that a flat-top kernel raised to a power can never become of “steep-origin” as it remains a flat-top; thus, the implied re-scaling will be unsuccessful, and flat-top kernels can not be used in this connection.

by:

$$\hat{f}_{jk} = \frac{1}{2\pi} \sum_{m=-T}^T \lambda_{g,c}(m/S_{jk}) \hat{\rho}_{jk}(m) e^{-imw}. \quad (11)$$

Because $\hat{\rho}_{jk}(m)$ is bounded (by unity), $Var(\hat{f}_{jk})$ is well-defined even if $\{V_t\}$ does not possess finite fourth moments. The following alternative to eq. (8) is then suggested:

$$Var(\hat{f}_{jk}) = O(S_{jk}/T) \text{ for any fixed } j, k. \quad (12)$$

Eq. (12) is now typically satisfied under regularity conditions; see e.g. Robinson (1991) and Hansen (1992) who considered the problem of spectral estimation in the absence of finite fourth moments.

A further consequence of lack of finite fourth moments is that, although $\hat{\rho}(m)$ will still be \sqrt{T} -consistent under appropriate weak dependence assumptions, $\hat{\Gamma}(m)$ is consistent but typically at slower rate; see e.g. Theorem 7.2.2 of Brockwell and Davis (1991) or Embrechts et al. (1997). A reasonable assumption adopted by Robinson (1991) is:

$$\hat{\Gamma}_{jj}(0) = \Gamma_{jj}(0) + O_P(1/T^\alpha), \text{ for all } j, \text{ and some } \alpha \in (0, 1/2]. \quad (13)$$

For our purposes we will require the slightly stronger condition:

$$E \left| \hat{\Gamma}_{jj}(0) - \Gamma_{jj}(0) \right|^{1+\delta} = O(1/T^{\alpha(1+\delta)}) \text{ for all } j, \text{ and some } \delta > 0 \text{ and } \alpha \in (0, 1/2]. \quad (14)$$

The following theorem is a generalization of Theorem 2.1 to the setting where finite fourth moments are potentially lacking.

Theorem 3.1 *Fix values for j, k , and assume conditions (12), (14), and that⁵*

$$S_{jk}^{-1} \sum_{j=-T+1}^{T-1} |\lambda_{g,c}(j/S_{jk})| = O(1). \quad (15)$$

Also assume $\Gamma_{jj}(0) > 0$ for all j .

(i) If $\sum_{m=-\infty}^{\infty} |m|^r |\Gamma_{jk}(m)| < \infty$ for some real number $r \geq 1$, then letting S_{jk} proportional to $T^{\alpha/(r+1)}$ yields

$$\hat{f}_{jk} = f_{jk} + O_P(T^{-\alpha r/(r+1)}), \quad (16)$$

⁵As in condition (i) of Lemma 8.1, eq. (15) is easily satisfied such as when $\lambda_{g,c}(x)$ has a hard ‘cut-off’, i.e., $\lambda_{g,c}(x) = 0$ for $|x| > \text{some } x_0$.

and

$$\hat{F}_{jk} = F_{jk} + O_P(T^{-\alpha r/(r+1)}). \quad (17)$$

(ii) If $|\Gamma_{jk}(m)| \leq Ce^{-am}$ for some constants $C, a > 0$, then letting $S_{jk} \sim A \log T$, for some appropriate constant A , yields

$$\hat{f}_{jk} = f_{jk} + O_P\left(\frac{\log T}{T^\alpha}\right) \quad \text{and} \quad \hat{F}_{jk} = F_{jk} + O_P\left(\frac{\log T}{T^\alpha}\right). \quad (18)$$

(iii) If $\Gamma_{jk}(m) = 0$ for $|m| > \text{some } q$, then letting $S_{jk} = \max(\lceil q/c \rceil, 1)$, yields

$$\hat{f}_{jk} = f_{jk} + O_P\left(\frac{\log \log T}{T^\alpha}\right) \quad \text{and} \quad \hat{F}_{jk} = F_{jk} + O_P\left(\frac{\log \log T}{T^\alpha}\right). \quad (19)$$

As a matter of fact, as apparent from the proof of part (iii), the $\log \log T$ term in the above eq. (19) could be replaced by a term that tends to infinity even slower, e.g. $\log \log \log T$.

So even under the potential absence of finite fourth moments, \hat{F} maintains its higher-order accuracy. Parts (ii) and (iii) of Theorem 3.1 show that the rate of convergence of \hat{F} comes very close to T^α which is the rate of convergence of $\hat{\Gamma}(0)$. Interestingly, under the premises of either part (ii) or (iii) of Theorem 3.1, the optimal rates for the bandwidth S_{jk} are insensitive to whether fourth moments are finite or not.

Qualitatively, the same is true for part (i). Note though that, as suggested by a referee, the rate achieved for \hat{F}_{jk} in part (i)—and the suggested choice for S_{jk} —does not appear to be sharp since it does not approach the corresponding optimal rate given in part (i) of Theorem 2.1 when $\alpha \rightarrow 1/2$. We conjecture, therefore, that the optimal rate of convergence of \hat{F}_{jk} under part (i) might be something like $O_P(T^{-\alpha r/(r+1/2)})$ perhaps times a logarithmic factor. Nevertheless, the difference of the above conjectured rate with the $O_P(T^{-\alpha r/(r+1)})$ rate given in Theorem 3.1 is small, and becomes negligible if r is large.

Remark 3.1 Until this point it was assumed, for simplicity, that the mean of our data is known and it is zero. We now show why this simplifying assumption can be made without loss of generality. To make this clear, here—and for the remainder of Section 3—assume that our observed data are the vectors X_1, \dots, X_T and that their mean $EX_t = \mu$ is unknown and estimated by $\bar{X}_T = T^{-1} \sum_{t=1}^T X_t$.

For $t = 1, \dots, T$, let $V_t = X_t - \mu$ and $\hat{V}_t = X_t - \bar{X}_T$, and define

$$\hat{\Gamma}(j) = \frac{1}{T} \sum_{t=1}^{T-j} \hat{V}_t \hat{V}'_{t+j} \quad \text{for } j \geq 0, \quad \text{and} \quad \hat{\Gamma}(j) = \hat{\Gamma}(-j)' \quad \text{for } j < 0,$$

letting $\hat{\Gamma}(j) = 0$ for $|j| \geq T$ as usual. Similarly, define the matrix $\hat{\rho}$ with (j, k) element given by $\hat{\rho}_{jk}(m) = \hat{\Gamma}_{jk}(m) / \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)}$. Finally, define

$$\hat{f}_{jk} = \frac{1}{2\pi} \sum_{m=-T}^T \lambda_{g,c}(m/S_{jk}) \hat{\rho}_{jk}(m) e^{-imw} \quad \text{and} \quad \hat{F}_{jk} = \frac{1}{2\pi} \sum_{m=-T}^T \lambda_{g,c}(m/S_{jk}) \hat{\Gamma}_{jk}(m) e^{-imw}.$$

The quantities $\hat{\Gamma}_j$, \hat{f} and \hat{F} are defined exactly as before as functions of V_1, \dots, V_T .

The key observation now is that $\hat{V}_t = V_t + M_T$ where $M_T \equiv \bar{X}_T - \mu = O_P(1/\sqrt{T})$ by Chebychev's inequality and the existence of the spectral density. Note also that

$$\hat{\Gamma}_{jk}(m) = \hat{\Gamma}_{jk}(m) - M_T^{(j)} M_T^{(k)} + M_T^{(j)} \frac{1}{T} \sum_{t=1}^{m-1} V_t^{(k)} + M_T^{(k)} \frac{1}{T} \sum_{t=m+1}^T V_t^{(j)}.$$

But the terms $\sum_{t=1}^{m-1} V_t^{(k)}$ and $\sum_{t=m+1}^T V_t^{(j)}$ are both of order $O_P(\sqrt{m})$. Since $\sqrt{m} \leq \sqrt{T}$, it follows that $\hat{\Gamma}(m) - \hat{\Gamma}(m) = O_P(1/T)$ uniformly in m , and the following corollary ensues.

Corollary 3.1 *Assume $\sum_{m=-\infty}^{\infty} |\Gamma_{jk}(m)| < \infty$ and eq. (15). Then,*

$$\hat{f}_{jk} = \hat{f}_{jk} + O_P(S_{jk}/T) \quad \text{and} \quad \hat{F}_{jk} = \hat{F}_{jk} + O_P(S_{jk}/T).$$

Hence, either part (i, ii or iii) of Theorem 2.1 or of Theorem 3.1 hold true as stated with \hat{f}_{jk} and \hat{F}_{jk} taking the place of f_{jk} and F_{jk} respectively.

The moral of Corollary 3.1 is that if the mean is unknown, plugging in the sample mean as an estimate has a negligible effect on the accuracy of spectral density estimators—provided, of course, that the bandwidth is chosen such that $S_{jk} = o(T)$. Ng and Perron (1996) give further discussion on the effect of an unknown mean in the spectral estimation set-up, and also provide exact, finite-sample expressions for the bias and variance of a spectral density estimator in the case of Gaussian noise.

Remark 3.2 Still in the case of Remark 3.1 with data X_1, \dots, X_T and unknown mean $EX_t = \mu$, it may be desirable to test the hypothesis $H_0 : \mu = \mu_0$. Assuming that $F(0)$ exists and is positive definite, the test statistic $\tau_{\mu_0} \equiv \sqrt{T} \hat{\Omega}^{-1/2} (\bar{X}_T - \mu_0)$ may be used where $\hat{\Omega}^{1/2}$ is a square root of matrix $\hat{\Omega} = 2\pi \hat{F}(0)$; of course, to define this square root and its inverse, the estimate $\hat{F}(0)$ must also be positive definite. The idea of studentizing a mean-like statistic by a nonparametric spectrum estimate can be traced back to Jowett (1954), Hannan (1957), and Brillinger (1979); see Robinson and Velasco (1997) for a review, and Velasco and Robinson (2001) for some recent developments.

Note that under any of the conditions—part (i), (ii) or (iii)—of Theorem 3.1, we can write $\hat{\Omega} = \Omega + O_P(1/T^\beta)$ for some $\beta \leq 1/2$. Since we are in the set-up of finite second moments, it follows that $\bar{X}_T = \mu + O_P(1/\sqrt{T})$, and thus—by the δ -method—we have

$$\tau_\mu \equiv \sqrt{T} \hat{\Omega}^{-1/2}(\bar{X}_T - \mu) = Z_T + O_P(1/T^\beta)$$

where $Z_T = \sqrt{T} \Omega^{-1/2}(\bar{X}_T - \mu)$. Under conditions validating the Central Limit Theorem, Z_T will tend to a multivariate standard normal distribution, and—by Slutsky’s theorem—so will τ_μ . For this reason the normal tables are used to approximate the appropriate quantile of the null distribution of τ_{μ_0} in order to obtain the threshold of the test.

This $O_P(1/T^\beta)$ discrepancy between τ_μ and Z_T is undesirable, and—if possible—should be minimized. To see why, under regularity conditions, a Berry-Esseen and/or Edgeworth expansion argument as discussed in Lahiri (2003) yields

$$\sup_z |Prob\{Z_T \leq z\} - \Phi(z)| = \begin{cases} O(1/\sqrt{T}) & \text{if } E|X_1|^3 < \infty \\ o(1/\sqrt{T}) & \text{if } EX_1^3 = 0 \\ 0 & \text{if the } X_t\text{s are jointly normal} \end{cases} \quad (20)$$

where $\Phi(z)$ is the distribution of $N(0,1)$. In other words, the distribution of Z_T will, in general, be quite close to standard normal; the same, however, does not hold for the distribution of τ_μ exactly due to the slow-converging $O_P(1/T^\beta)$ discrepancy. Furthermore, due to the addition of this $O_P(1/T^\beta)$ term, the variance of any of the coordinates of τ_μ will, in finite-samples, tend to be larger than the variance of the corresponding coordinate of Z_T .

So, both because of the slow convergence to normality as well as the higher variance, the standard normal tables will tend to give suboptimal approximations to the quantiles of τ_μ (and of τ_{μ_0} under the null). This phenomenon is manifested even when the X_t s are independent (and Gaussian) in which case the $O_P(1/T^\beta)$ term (with $\beta = 1/2$) is responsible for the difference between the z -tables and the t -tables. In the time series case where typically $\beta < 1/2$ the effects of this phenomenon are quite more pronounced, and are summarized as:

- **Size.** The test threshold obtained from the z -tables will be smaller (in absolute value) than the ideal; the implication is that the size of the test will be larger than nominal.
- **Power.** The power of the test is compromised because of the higher variance associated with the $O_P(1/T^\beta)$ term.⁶

⁶Interestingly, the inaccurate choice of too small a threshold mentioned above works in the opposite

The moral is that the $O_P(1/T^\beta)$ discrepancy can be minimized—by maximizing β —, and here is where the flat-top kernels may be useful. Recall that $\beta = 1/3$ for the Newey-West (1987) estimator whereas $\beta = 2/5$ for all second order kernels considered by Andrews (1991). By contrast, under any of the conditions of Theorem 2.1, the flat-top kernels will achieve a higher value for β which—in parts (ii) and (iii) of the theorem—becomes practically equal to $1/2$. It turns out from the lower bounds of Samarov (1977) that the β achieved by the flat-top kernels is actually optimal under the three conditions of Theorem 2.1. Although there is no literature on the optimal rates for spectral estimators in a setting where finite fourth moments are lacking, we conjecture that the rate of convergence achieved by the flat-top kernels under conditions (ii) or (iii) of Theorem 3.1 is optimal, whereas—as previously mentioned—the rate under condition (i) is very close to optimal.

As discussed in the Introduction, the flat-top kernels are not positive semi-definite, and hence their use for studentization purposes has been limited to date since there is no guarantee that $\hat{\Omega}^{-1/2}$ is well-defined. It should be stressed, however, that if the estimand $F(0)$ is positive definite, then $\hat{F}(0)$ will be positive definite with probability that quickly tends to one; see also Note 4 of Andrews (1992), and Remark 4.1 in what follows. The above discussion is corroborated by our finite-sample simulations of Section 7; in particular, in models where the spectral density matrix was well-conditioned (e.g. Models I and III of Section 7.1), the proportion of occurrences of non-positive semi-definite estimated matrices was found to be very small for a sample size of $T = 100$, and it was practically *zero* (of order $1/10,000$) for $T = 500$.

In any case, when an absolute safeguard against non-positive definiteness is needed, i.e., when a practitioner needs the positive definiteness to hold with probability *exactly* one, a simple modification of \hat{F} is readily available. In particular, it is shown in the next section how a positive semi-definite (or even definite) estimated spectral density matrix can be obtained *without* compromising the high rate of convergence associated with flat-top kernels. This type of fast-converging albeit positive definite spectral (or HAC covariance) matrix can find immediate application in set-ups like the studentized bootstrap; see Lahiri (2003), or Inoue and Shintani (2006).

direction and may have a partially off-setting effect.

4 Positive semi-definite spectral estimation

Flat-top kernels are infinite-order kernels, and therefore they are capable of achieving higher-order accuracy when that is possible. For example, it is apparent that, under the MA(q)-type condition of Theorem 2.1 (iii), \sqrt{T} -consistent estimation of F_{jk} is possible since F_{jk} is a function of only finitely many (q) parameters. The flat-top estimator \hat{F}_{jk} indeed attains \sqrt{T} -consistency in that case, and the flatness of the kernel over the interval $[-c, c]$ is crucial for this attainment.

The only disadvantage of flat-top kernels is that they are not positive semi-definite, i.e., the matrix \hat{F} is not almost surely positive semi-definite for all w . Nevertheless, the fast rate of convergence of \hat{F} to a positive semi-definite matrix indicates that the incidents of a non-positive semi-definite \hat{F} may be rare; this fact was documented in the simulations of Andrews (1991) with respect to the truncated kernel that technically belongs to the flat-top family.⁷

However, the positive semi-definiteness is an important property especially in the case of $w = 0$ when the object is estimation of a covariance matrix; see e.g. the testing example of Remark 3.2 where, in particular, positive *definiteness* is required. It is likely for this reason that the focus in the recent literature starting with Newey-West (1987) has been on positive semi-definite estimators. Nonetheless, we now show how the flat-top estimator \hat{F} can be easily modified to render a positive semi-definite estimator.

Recall that a Hermitian matrix has all real eigenvalues, and can be diagonalized by a unitary transformation. Thus, consider the unitary decompositions of the Hermitian matrices F and \hat{F} , namely:

$$F = U\Lambda U^* \quad \text{and} \quad \hat{F} = \hat{U}\hat{\Lambda}\hat{U}^* \tag{21}$$

where U, \hat{U} are unitary (complex-valued) matrices, i.e., they satisfy $U^{-1} = U^*$ and $\hat{U}^{-1} = \hat{U}^*$ where $*$ denotes the conjugate transpose; the columns of U and \hat{U} are the orthonormal eigenvectors of F and \hat{F} respectively, and $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_d)$, $\hat{\Lambda} = \text{diag}(\hat{\lambda}_1, \dots, \hat{\lambda}_d)$ are

⁷Note, however, that the discontinuity of the truncated kernel gives its corresponding spectral window very pronounced ‘sidelobes’, and hence high variance (because of large l_2 -norm) and unfavorable finite-sample behavior that have been widely reported; see Figure 1 in Politis and Romano (1995) for a comparative graph of the sidelobes. Because of its discontinuity, the truncated kernel is arguably the *worst* representative of the flat-top family; more details on kernel choice are given in Section 5.

diagonal matrices containing the respective eigenvalues.

Noting that the entries of $\hat{\Lambda}$ are all nonnegative suggests the following fix to the possible negativity of \hat{F} . Let $\hat{\Lambda}^+ = \text{diag}(\hat{\lambda}_1^+, \dots, \hat{\lambda}_d^+)$ where $\hat{\lambda}_j^+ = \max(\hat{\lambda}_j, 0)$, i.e., the entries of $\hat{\Lambda}^+$ are given by the positive part of the entries of $\hat{\Lambda}$, and define the positive semi-definite estimator

$$\hat{F}^+ = \hat{U}\hat{\Lambda}^+\hat{U}^*. \quad (22)$$

The following theorem shows that, in addition to being positive semi-definite, \hat{F}^+ inherits the higher-order accuracy of \hat{F} . Therefore, \hat{F}^+ is our proposed higher-order accurate, positive semi-definite estimator.

Theorem 4.1 *Let R_T be a sequence such that $R_T \rightarrow \infty$ as $T \rightarrow \infty$. If $\hat{F} = F + O_P(1/R_T)$, then $\hat{F}^+ = F + O_P(1/R_T)$ as well.⁸*

To take it one step further, it may be the case that the estimand F is not only positive semi-definite but strictly positive definite. Alternatively, it can be of interest to consider the inverse of an estimated F as in the case of the studentized statistic⁹ of Remark 3.2. For such applications, it may be desirable to have a *strictly* positive definite estimator of F that maintains the high accuracy of the flat-top estimators. For this reason, let $\epsilon_T > 0$ be some chosen sequence, and define $\hat{\Lambda}^\epsilon = \text{diag}(\hat{\lambda}_1^\epsilon, \dots, \hat{\lambda}_d^\epsilon)$ where $\hat{\lambda}_j^\epsilon = \max(\hat{\lambda}_j, \epsilon_T)$. Also define the strictly positive definite estimator

$$\hat{F}^\epsilon = \hat{U}\hat{\Lambda}^\epsilon\hat{U}^*. \quad (23)$$

The following corollary to Theorem 4.1 shows that \hat{F}^ϵ also inherits the higher-order accuracy of \hat{F} if ϵ_T is chosen right. Thus, \hat{F}^ϵ is a higher-order accurate, strictly positive definite estimator.

Corollary 4.1 *Let R_T be a sequence such that $R_T \rightarrow \infty$ as $T \rightarrow \infty$, and let the strictly positive sequence ϵ_T be $o(1/R_T)$. If $\hat{F} = F + O_P(1/R_T)$, then $\hat{F}^\epsilon = F + O_P(1/R_T)$ as well.*

⁸The notation $A = O_P(1/R_T)$ for some matrix A means that each element of A is $O_P(1/R_T)$.

⁹Note, however, that it may be possible to conduct a hypothesis test even with an estimated covariance matrix that is just positive *semi*-definite; see the discussion in Section 7.2.

Note that for spectral estimation problems we always have $1/\sqrt{T} = O(1/R_T)$. So any choice of $\epsilon_T > 0$ satisfying $\epsilon_T = o(1/\sqrt{T})$ will satisfy the requirements of Corollary 4.1. However, in order to avoid the introduction of unnecessary finite-sample bias it is recommended to take $\epsilon_T > 0$ quite smaller than $1/\sqrt{T}$. But then again ϵ_T should not be too small in order not to risk the matrix \hat{F}^ϵ being ill-conditioned which leads to computational difficulties; letting $\epsilon_T = 1/T^a$ with $1 \leq a \leq 2$ seems like a reasonable practical compromise.¹⁰

Remark 4.1 Some concluding remarks are in order here. Since Theorems 2.1 and 3.1 show that the flat-top estimator \hat{F} is consistent at a very fast rate, we expect it to be close to its target value F and share its properties (positive definiteness, etc.). This is indeed true, and supported by the finite-sample simulations of Section 7.

To elaborate, if the eigenvalues of the estimand F are not close to zero, then with high probability the eigenvalues of \hat{F} will be positive as well and there is no need for \hat{F}^+ or even \hat{F}^ϵ . As a matter of fact, from eq. (50) it follows that the probability \hat{F} is *not* positive definite—when F is—is of order $O(\tilde{S}/T)$ where $\tilde{S} = \max_{j,k} S_{jk}$. Hence, in the set-up of the ARMA-type condition of part (ii) of Theorem 2.1, the probability that \hat{F} will be positive definite is of order $1 - O(\log T/T)$.

On the other hand, if an eigenvalue of F is zero (or close to zero), then the small bias of \hat{F} *demands* that the corresponding eigenvalue of \hat{F} has a distribution that is centered right around zero which, as a consequence, generates many negative values (as many as 50%); see Figure 3 (b) for an illustration. However, this is *not* to be seen as a hindrance; rather, it is very informative, giving strong evidence that the target eigenvalue is close to zero, and that consequently adopting \hat{F}^+ or \hat{F}^ϵ is appropriate, resulting in a highly-accurate estimator.

Consider for example the one-dimensional case ($d = 1$), and note that the asymptotic normality of kernel estimators discussed in Remark 2.2 implies that the desire for positivity clashes with the desire for unbiasedness; this is especially apparent either in small/medium-size samples, or in large samples with a target value near zero. In other words, restricting our attention to just non-negative estimators is tantamount to limiting ourselves to working with severely biased estimators; see e.g. Figure 3 (a).

¹⁰The estimator \hat{F}^ϵ as defined above is not scale equivariant since ϵ_T is not; see Remark 6.1 for a discussion. A way to remedy this deficiency is to work with the normalized spectral density matrix f instead. I.e., if $\hat{\xi}_j$ is the j th eigenvalue of \hat{f} , define \hat{f}^ϵ as having eigenvalues $\hat{\xi}_j^\epsilon = \max(\hat{\xi}_j, \epsilon_T)$ for $j = 1, 2, \dots$ and the same eigenvectors as \hat{f} ; finally, obtain \hat{F}^ϵ from \hat{f}^ϵ by multiplying with the appropriate (estimated) variance/covariance element-by-element.

The thesis of this paper is to not impose the non-negativity restriction at the outset; rather, to work with the most accurate (less biased) estimators, and fix the possible non-positivities at the end. Because of the high accuracy of the proposed estimators, non-positivities will be observed in practice effectively only when the target value is zero (or close to zero) in which case an estimated value of zero (or positive but close to zero) is right on target. The finite-sample simulations of Section 7 corroborate this claim.

5 Flat-top kernel choice

The favorable asymptotic rates of Theorems 2.1 and 3.1 are achievable by any member of the flat-top family. Nevertheless, finite-sample properties will be dependent upon kernel choice. For example, as mentioned in the previous section, the truncated kernel $\kappa_{trunc}(x)$ is one of the worse representatives of the flat-top family because of the pronounced ‘sidelobes’ of the Dirichlet kernel which is its corresponding spectral window—see e.g. Figure 2 of Politis and Romano (1995). Since half of those sidelobes are on the negative side, they unnecessarily inflate the L_2 -norm of the spectral window under the constraint that the latter integrates to one; as implied by Remark 2.2, a large L_2 -norm results in a large variance.¹¹

In order to reduce the size of a spectral window’s sidelobes, the flat-top kernel must be chosen as smooth as possible. The poor finite-sample performance of the truncated kernel is actually due to the discontinuity of the function $\kappa_{trunc}(x)$ at points ± 1 . The trapezoidal kernel $\lambda_{TR,c}(x)$ is continuous everywhere, and is thus much better performing than the truncated. Even better finite-sample behavior is expected if the ‘corners’ of the trapezoid $\lambda_{TR,c}(x)$ are smoothed out. For example, McMurry and Politis (2004) constructed a member of the flat-top family that is infinitely differentiable; it is defined as

$$\lambda_{ID,b,c}(x) = \begin{cases} 1 & \text{if } |x| \leq c \\ \exp(-b \exp(-b/(|x| - c)^2)/(|x| - 1)^2) & \text{if } c < |x| < 1 \\ 0 & \text{if } |x| \geq 1 \end{cases} \quad (24)$$

where $c \in (0, 1]$, and $b > 0$ is a shape parameter.

Nevertheless, the already good performance of the trapezoidal kernel indicates that one might not have to use an infinitely differentiable kernel to gather appreciable finite-sample

¹¹The variance is still of order $O(S_{jk}/T)$ as eq. (8) demands, but the proportionality constant in the term $O(S_{jk}/T)$ is large for the Dirichlet kernel.

benefits. For example, we can create a flat-top kernel by adding a piecewise cubic tail, similar to that of Parzen's (1961) kernel, to the $[-c, c]$ flat-top region. The resulting flat-top kernel would be defined as:

$$\lambda_{PR,c}(x) = \begin{cases} 1 & \text{if } 0 \leq x \leq c \\ 1 - 6(x - c)^2 + 6|x - c|^3 & \text{if } c \leq x \leq c + 1/2 \\ 2(1 - |x - c|)^3 & \text{if } c + 1/2 < x < c + 1 \\ 0 & \text{if } x \geq c + 1 \\ \lambda_{PR,c}(-x) & \text{if } x < 0. \end{cases} \quad (25)$$

The original Parzen kernel $\kappa_{PR}(x)$ is seen to be equal to $\lambda_{PR,0}(x)$. Similarly, we can create a flat-top kernel by a modification of Priestley's (1962) 'quadratic spectral kernel':

$$\kappa_{QS}(x) = \frac{25}{12\pi^2 x^2} \left(\frac{\sin(6\pi x/5)}{6\pi x/5} - \cos(6\pi x/5) \right)$$

that has been found optimal¹² among positive semi-definite second order kernels; see e.g. Priestley (1962) or Epanechnikov (1969). The modification would amount to defining:

$$\lambda_{QS,b,c}(x) = \begin{cases} 1 & \text{if } 0 \leq x \leq c \\ \frac{3}{b^2(x-c)^2} \left(\frac{\sin(b(x-c))}{b(x-c)} - \cos(b(x-c)) \right) & \text{if } x > c \\ \lambda_{QS,b,c}(-x) & \text{if } x < 0, \end{cases} \quad (26)$$

so that $\lambda_{QS,b,c}(x)$ has the required $[-c, c]$ flat-top region, but inherits the tails of $\kappa_{QS}(x)$. Note that $\kappa_{QS}(x)$ tends to zero for large x but does not vanish after a cut-off point. The parameter $b > 0$ in $\lambda_{QS,b,c}(x)$ is again a shape parameter scaling the magnitude of the tail. Since c 'scales' together with b , we can let $c = 1$ in connection with $\lambda_{QS,b,c}(x)$, so that b is the only remaining shape parameter.

Having chosen the shape of the function g , the remaining parameters c and/or b have to be chosen as well. For the trapezoidal kernel $\lambda_{TR,c}(x)$, the recommendation of Politis and Romano (1995) is to take c in the neighborhood of $1/2$; the rationale is that the extreme values $c \rightarrow 0$ and $c \rightarrow 1$ are both to be avoided, corresponding to the aforementioned poorly performing kernels, the Bartlett and truncated kernel respectively.

¹²Priestley's kernel $\kappa_{QS}(x)$ leads to the so-called Epanechnikov spectral window of quadratic form, i.e., $K_{QS}(w) = (1 - w^2)^+$ that satisfies a number of optimality criteria among positive semi-definite second order kernels; see Andrews (1991).

For the infinitely differentiable kernel $\lambda_{ID,b,c}(x)$ there is an interplay between the two parameters b and c ; for example, even with c close to 0, there is a range of values of b that will make $\lambda_{ID,b,c}(x)$ look very much like the trapezoidal $\lambda_{TR,1/2}(x)$ with ultra-smoothed corners. Similarly, to implement the kernels $\lambda_{PR,c}(x)$ and/or $\lambda_{QS,b,1}(x)$, the parameters c and b must be chosen respectively.

The problem of identifying the optimal shape of a flat-top kernel is still open, and more work is needed in that respect. In the meantime, motivated by the good performance of the trapezoidal kernel $\lambda_{TR,1/2}(x)$, the following rule-of-thumb may be suggested: choose the parameter(s) of a flat-top kernel such that the resulting shape is similar to $\lambda_{TR,1/2}(x)$ with smoothed corners. For example, letting $c = 0.05$ and $b = 1/4$ has this desired effect in connection with $\lambda_{ID,b,c}(x)$, i.e., $\lambda_{ID,0.25,0.05}(x)$ ‘looks’ like a smoothed version of $\lambda_{TR,1/2}(x)$. To get $\lambda_{PR,c}(x)$ and $\lambda_{QS,b,1}(x)$ to yield a similar balance between the flat-top region and the tail, the values $c = 0.75$ and $b = 4$ may be used respectively. Plots of the flat-top kernels $\lambda_{TR,1/2}(x)$, $\lambda_{ID,0.25,0.05}(x)$, $\lambda_{PR,0.75}(x)$ and $\lambda_{QS,4,1}(x)$ are shown in Figure 1.

6 Adaptive bandwidth choice

In this section, assume that a member of the flat-top family, say $\lambda_{g,c}$, has been identified to be used for \hat{F}^+ and \hat{F} . Besides the favorable asymptotic properties and speed of convergence associated with flat-top kernels as demonstrated in Theorems 2.1 and 3.1, a further reason for using a flat-top lag-window is that choosing its bandwidth in practice is intuitive and doable by a simple inspection of the correlogram/cross-correlogram, i.e., a plot of $\hat{\rho}_{jk}(m)$ vs. m where $\hat{\rho}_{jk}(m) = \hat{\Gamma}_{jk}(m) / \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)}$ for all j, k .

The proposed bandwidth choice rule is motivated by case (iii) of Theorems 2.1 and 3.1 and boils down to looking for a point, say \hat{q} , after which the correlogram appears negligible, i.e., $\hat{\rho}_{jk}(m) \simeq 0$ for $|m| > \hat{q}$ (but $\hat{\rho}_{jk}(\hat{q}) \neq 0$). Of course, $\hat{\rho}_{jk}(m) \simeq 0$ is taken to mean that $\hat{\rho}_{jk}(m)$ is not significantly different from zero, i.e., an implied hypothesis test. After identifying \hat{q} , the recommendation is to just take $\hat{S}_{jk} = \max([\hat{q}/c], 1)$ as part (iii) of Theorems 2.1 and 3.1 suggests. Although it may be overoptimistic to expect that our data will follow a finite-order MA(q) model, the validity of this simple rule in general situations is due to the fact that an MA(q) model—with high enough q —can always serve as an approximation at least as far as the spectral density is concerned; see e.g. Brockwell and Davis (1991).




Figure 1: (a) Plot of $\lambda_{TR,1/2}(x)$ vs. $x > 0$; (b) Plot of $\lambda_{ID,0.25,0.05}(x)$ vs. $x > 0$; (c) Plot of $\lambda_{PR,0.75}(x)$ vs. $x > 0$; (d) Plot of $\lambda_{QS,4,1}(x)$ vs. $x > 0$.

The intuitive interpretation of the above bandwidth choice rule is an effort to extend the ‘flat-top’ region of $\lambda_{g,c}$ over the whole of the region where $\hat{\rho}_{jk}(m)$ is thought to be significant so as not to downweigh it and introduce bias. Nevertheless, the ‘flat-top’ region of $\lambda_{g,c}$ can be greater than $[-c, c]$ depending on the choice of function g . Even if $g(x)$ is strictly decreasing for $x > c$, its rate of decrease near c may be slow enough so that $\lambda_{g,c}(x) \simeq 1$ for x in an interval much greater than $[-c, c]$; see, for example, Figure 1 (b) regarding the infinitely differentiable $\lambda_{IS,b,c}(s)$ with $b = 1/4$ and $c = 0.05$. Thus, we are led to define the ‘effective’ flat-top region of $\lambda_{g,c}$ as the interval $[-c_{ef}, c_{ef}]$ where c_{ef} is the largest number such that $\lambda_{g,c}(x) \geq 1 - \epsilon$ for all x in $[-c_{ef}, c_{ef}]$; here ϵ is some small chosen number, e.g. $\epsilon = 0.01$.

Now we can rigorously define the empirical bandwidth choice rule. Note that in the case $j \neq k$, $\rho_{jk}(m)$ is the cross-correlation sequence which is not symmetric in m ; rather than looking at both positive and negative m , we choose to look at both $\rho_{jk}(m)$ and $\rho_{kj}(m)$ for only positive m which is equivalent.

EMPIRICAL RULE OF CHOOSING S_{jk} FOR FLAT-TOP KERNEL $\lambda_{g,c}$.

Case $j = k$: Let \hat{q} be the smallest nonnegative integer such that $|\hat{\rho}_{jk}(\hat{q}+m)| < C_0\sqrt{\log_{10} T/T}$, for $m = 0, 1, \dots, K_T$, where $C_0 > 0$ is a fixed constant, and K_T is a positive, nondecreasing integer-valued function of T such that $K_T = o(\log T)$. Then, let $\hat{S}_{jk} = \max(\lceil \hat{q}/c_{ef} \rceil, 1)$.

Case $j \neq k$: Let \hat{q}_{jk} be the smallest nonnegative integer such that $|\hat{\rho}_{jk}(\hat{q}_{jk} + m)| < C_0\sqrt{\log_{10} T/T}$, for $m = 0, 1, \dots, K_T$, where $C_0 > 0$ is a fixed constant, and K_T is a positive, nondecreasing integer-valued function of T such that $K_T = o(\log T)$. Similarly, let \hat{q}_{kj} be the smallest nonnegative integer such that $|\hat{\rho}_{kj}(\hat{q}_{kj} + m)| < C_0\sqrt{\log_{10} T/T}$, for $m = 0, 1, \dots, K_T$. Then, let $\hat{q} = \max(\hat{q}_{jk}, \hat{q}_{kj})$, and $\hat{S}_{jk} = \hat{S}_{kj} = \max(\lceil \hat{q}/c_{ef} \rceil, 1)$.

In the univariate case (i.e., $d = 1$ or $j = k$ in the above), the bandwidth choice rule was empirically suggested by Politis and Romano (1995) for the trapezoidal kernel; it was then rigorously studied in Politis (2003) but still only for the trapezoidal kernel. Note that the constant C_0 and the form of K_T are the practitioner’s choice. Politis (2003) makes the concrete recommendations $C_0 \simeq 2$ and $K_T = \max(5, \sqrt{\log_{10} T})$ that have the interpretation of yielding (approximate) 95% simultaneous confidence intervals for $\rho_{jk}(\hat{q}+m)$ with $m = 1, \dots, K_T$ by Bonferroni’s inequality.

These approximate confidence intervals are based on a null hypothesis that the series is

i.i.d. in which case the large-sample variance of $\rho_{jk}(m)$ is $1/T$ for any m . Nevertheless, in nonlinear/non-normal time series, uncorrelatedness is a weaker assumption than independence. So, a more conservative approach would be to use a resampling and/or subsampling approach—cf. Lahiri (2003) or Politis, Romano and Wolf (1999)—in order to estimate these variances and adjust C_0 appropriately. For example, if the standard deviation of $\sqrt{T}\rho_{jk}(m)$ for some m among the lags under consideration is estimated by ν , then let $C_0 \simeq 2\nu$ and $K_T = \max(5, \sqrt{\log_{10} T})$ as before.

In any case, the practitioner should always be vigilant in a case where altering the value of C_0 slightly leads to radically different values of \hat{q} . In such a case, the rule-of-thumb is to use the smaller of the two potential estimates \hat{q} in the sense that flat-top kernels work best with small bandwidth parameters; see Politis and White (2004) for an example of this phenomenon.

Remark 6.1 As pointed out by a referee, a disadvantage of element-by-element bandwidth choice methods is that they yield spectral estimators that are not equivariant. When Q is a full-rank matrix and the data is transformed from V_t to QV_t , an equivariant estimator is transformed from \hat{F} to $Q\hat{F}Q'$. Nevertheless, the above empirical rule for choosing the bandwidths is, at least, equivariant with respect to *scaling*, i.e., when the matrix Q is diagonal.

The performance of our empirical bandwidth choice rule is quantified in the following theorem. The case $j = k$ of the theorem was first given in Politis (2003) for the trapezoidal flat-top kernel; a similar theorem is also given in Berg and Politis (2008). It is important to note that the theorem is valid even under the setting of Section 3, i.e., under a potential lack of finite fourth moments.

Theorem 6.1 *Fix j, k , and assume conditions strong enough to ensure that¹³ for all finite N ,*

$$\max_{m=1, \dots, N} |\hat{\rho}_{jk}(n+m) - \rho_{jk}(n+m)| = O_P(1/\sqrt{T}) \quad (27)$$

¹³There exist different sets of conditions sufficient for eq. (27); see Brockwell and Davis (1991) or Romano and Thombs (1996). As a matter of fact, under further regularity conditions, the process $\sqrt{T}(\hat{\rho}_{jk}(\cdot) - \rho_{jk}(\cdot))$ is asymptotically Gaussian with autocovariance tending to zero; consequently, eq. (28) would follow from the theory of extremes of dependent sequences—see e.g. Leadbetter et al. (1983) or Ch. 5 of Hannan and Deistler (1988).

uniformly in n , and

$$\max_{m=0,1,\dots,T-1} |\hat{\rho}_{jk}(m) - \rho_{jk}(m)| = O_P\left(\sqrt{\frac{\log T}{T}}\right). \quad (28)$$

Also assume that the sequence $\rho_{jk}(m)$ does not have more than $K_T - 1$ consecutive zeros¹⁴ in its first m_0 lags (i.e., for $m = 0, 1, \dots, m_0$).

(i) Assume that for $m > m_0$ we have $\rho_{jk}(m) = C_1 m^{-p_1}$ or $\rho_{jk}(m) = C_1 m^{-p_1} \cos(a_1 m + \theta_1)$, and $\rho_{kj}(m) = C_2 m^{-p_2}$ or $\rho_{kj}(m) = C_2 m^{-p_2} \cos(a_2 m + \theta_2)$, for some positive integers p_1, p_2 , and some constants satisfying $C_v > 0$, $a_v \geq \frac{\pi}{K_T}$, and $\theta_v \in [0, 2\pi]$ for $v = 1, 2$. Then,

$$\hat{S}_{jk} \stackrel{P}{\sim} \frac{A_1 T^{1/(2p)}}{(\log T)^{1/(2p)}} \quad \text{where } p = \max(p_1, p_2)$$

for some positive constant A_1 ; the notation $A \stackrel{P}{\sim} B$ means $A/B \xrightarrow{P} 1$.

(ii) Assume that for $m > m_0$ we have $\rho_{jk}(m) = C_1 \xi_1^m$ or $\rho_{jk}(m) = C_1 \xi_1^m \cos(a_1 m + \theta_1)$, and $\rho_{kj}(m) = C_2 \xi_2^m$ or $\rho_{kj}(m) = C_2 \xi_2^m \cos(a_2 m + \theta_2)$, where the constants satisfy $C_v > 0$, $|\xi_v| < 1$, $a_v \geq \frac{\pi}{K_T}$, and $\theta_v \in [0, 2\pi]$ for $v = 1, 2$. Then $\hat{S}_{jk} \stackrel{P}{\sim} A_2 \log T$ where $A_2 = -1/\max(\log |\xi_1|, \log |\xi_2|)$.

(iii) If $|\rho_{jk}(m)| + |\rho_{kj}(m)| = 0$ for $m > \text{some nonnegative integer } q$ (with $q < m_0 + K_T$), but $|\rho_{jk}(q)| + |\rho_{kj}(q)| \neq 0$, then $\hat{S}_{jk} = \max(\lceil q/c_{ef} \rceil, 1) + o_P(1)$.

Comparing the empirical rule \hat{S}_{jk} to the theoretically optimal values of S_{jk} given in Theorem 2.1 we see that \hat{S}_{jk} manages to capture exactly the theoretically optimal rate in cases (ii) and (iii) of Theorem 6.1. In case (i) of Theorem 6.1, \hat{S}_{jk} increases essentially as a power of T since the $2p$ -th root of the logarithm changes in an ultra-slow way with T ; note that the empirically found exponent $1/(2p)$ is slightly smaller than the theoretically optimal bandwidth given in part (i) of Theorem 2.1 but the difference is small, and becomes even smaller for large p .

Thus, \hat{S}_{jk} is seen to *adapt* to the underlying rate of decay of the correlation and cross-correlation functions, automatically switching between the polynomial, logarithmic, and constant rates that are optimal respectively in the three cases of Theorem 2.1. In that sense, the literature on adaptive nonparametric estimation such as Lepski (1990, 1991, 1992) may be quite relevant. In particular, note that our empirical bandwidth choice rule

¹⁴Because of this assumption, it is advisable to take K_T be an increasing function of T , albeit at the very slow rate suggested by the recommendation $K_T = \max(5, \sqrt{\log_{10} T})$.

has the qualitative flavor of wavelet ‘thresholding’ of Donoho and Johnstone (1995) albeit remaining in the context of linear estimators.

Remark 6.2 As already made clear, the proposed empirical bandwidth choice rule only attempts (and, to a large extent, succeeds) to capture the correct *rates* under the three autocovariance conditions of Theorems 2.1 or 3.1; it does not necessarily give the optimal constants associated with these rates. To elaborate, the usual plug-in methods—such as Andrews’s (1991)—focus on estimating (and plugging-in) the constants in the leading terms of bias and variance of \hat{F} with a subsequent MSE minimization to get the estimated bandwidth. However, these methods can not adapt to an arbitrary degree of smoothness of the underlying function; their saving point is that the leading term for the bias is dictated by the order of the kernel. By contrast, there is no such limitation in the flat-top setting since the order is infinite; the implication is that (a) the order of magnitude of the bias is not dictated by the kernel but by the underlying smoothness which must, therefore, be identified, and (b) the leading term of the bias is elusive except for an upper bound. For both of these reasons, the focus of the empirical bandwidth choice rule is the (implicit) identification of which of the three autocovariance conditions holds for the data at hand, and not the estimation of optimal constants.

Focusing on the autocorrelation—as the empirical rule does—has an additional, very important by-product: Theorem 6.1 is valid as stated *even* under a potential lack of finite fourth moments. Notably, the aforementioned plug-in methods fail in such a case since the variance of \hat{F} is infinite. So, even without finite fourth moments, \hat{S}_{jk} adapts to the underlying rate of decay of the correlation and cross-correlation functions, automatically switching between the polynomial, logarithmic, and constant rates that are optimal respectively in the three cases of Theorem 3.1 as well.¹⁵

7 Some finite-sample simulations

We now present some finite-sample simulations to complement our asymptotic results. The simulations are not meant to be exhaustive; rather, their goal is to illustrate the main issues discussed in the paper. Throughout Section 7, all simulations were done in the realistic

¹⁵Under the caveat that in the case (i) of Theorem 3.1 the optimal polynomial rate for S_{jk} is not known; see the discussion after Theorem 3.1.

setting of an unknown mean, and the data were centered by the sample mean to estimate the covariances. However, in this section, the simple notation \hat{F} was used to denote our estimated matrices, rather than the more cumbersome $\hat{\hat{F}}$ notation used in Remark 3.1.

The bandwidths of the ‘traditional’ kernels κ_B (Bartlett), κ_{PR} (Parzen), and κ_{QS} (optimal 2nd order kernel) were estimated using equations (6.2) and (6.4) of Andrews (1991), i.e., the notion of estimating the bandwidth constants by fitting an AR(1) model. By contrast, the bandwidths of all flat-top kernels were estimated using our empirical rule of Section 6. For the truncated kernel κ_{trunc} both bandwidth choices, i.e., the Andrews bandwidth—see footnote 5 in Andrews (1991, p. 834)—and our empirical rule, were used and are denoted by Truncated-A and Truncated-E respectively.

7.1 Estimating $F(0)$

We now focus on estimating $F(w)$ with $w = 0$ for bivariate series ($d = 2$) generated by three simple models. For the simulation, $B = 999$ bivariate time series stretches, each of length T , were generated; the first two models were:

MODEL I: $V_t^{(1)} = 0.75V_{t-1}^{(1)} + Z_t^{(1)}$, and $V_t^{(2)} = 2(Z_t^{(2)} + Z_{t-1}^{(2)})$ where $V_t^{(k)}$ denotes the k th coordinate series of the bivariate series $\{V_t\}$.

MODEL II: $V_t^{(1)} = Z_t^{(1)} - Z_{t-1}^{(1)}$, and $V_t^{(2)} = W_t + V_{t+7}^{(1)}$ where $W_t = -0.75W_{t-1} + Z_t^{(2)}$.

In all the above, the error series $\{Z_t^{(1)}\}$ and $\{Z_t^{(2)}\}$ are i.i.d. standard normal and independent to each other.

Model I involves two coordinates independent to each other, an AR(1) and a MA(1), both exhibiting positive dependence. The independence of the two coordinates implies that $F_{12}(w) = 0$ for all w which in turn implies that the optimal value of S_{12} for the flat-top kernels is as small as possible, i.e, one; the other target values are $F_{11}(0) = 8/\pi = F_{22}(0)$.

Table 1a shows the empirically found MSEs of different estimators relative to (i.e., divided by) the MSE of the optimal second order estimator with kernel κ_{QS} ; the data followed Model I with $T = 100$. It is apparent that in the case of F_{11} all traditional kernels (Bartlett, Parzen and the optimal κ_{QS}) do quite well and outperform the recommended flat-top kernels of Figure 1. However, this seems to be due to the fact that we are using an AR(1) formula for the bandwidths of traditional kernels and an AR(1) model happens to

be correct in this case. That the bandwidth is the most prominent issue here is manifested by comparing the truncated kernel with AR(1) bandwidth (Truncated-A) to the one with bandwidth estimated by our empirical rule (Truncated-E). In fact, Truncated-A seems to be the overall *best* estimator of the AR(1) spectrum F_{11} with strong positive dependence present; this is consistent with the findings in Table II of Andrews (1991).

The situation is reversed in the estimation of F_{12} and F_{22} . Here, the problematic use of the same bandwidth for all coordinates of the target matrix F is apparent as Truncated-E, having coordinate-specific estimated bandwidth, outperforms Truncated-A. In fact, the best estimator of F_{12} and F_{22} appears to be the positive semi-definite estimator \hat{F}^+ corresponding to the truncated kernel with bandwidth matrix estimated by our empirical rule (Truncated-E).

This may not seem surprising since *had we known* that an MA(1) model holds for F_{22} , we would estimate F_{22} by a model-based estimator that would be tantamount to a truncated estimator in this case. However, note that the MA(1) information is *not* used here; rather, our empirical rule is able to sense and automatically adapt to this MA(1) structure, and this is a major success with a sample size as small as 100.

Figure 2 (a) shows a histogram of our empirical rule \hat{S}_{11} for use with the trapezoidal kernel $\lambda_{TR,1/2}$ as computed over the 999 Monte Carlo iterations. The mean of the histogram is about 9 which is right about what we would use had we known that the underlying model is an AR(1). It is the variability in this histogram that inflates the variances of our flat-top estimators with estimated bandwidths.

A histogram of the corresponding \hat{S}_{22} is not very informative as the overwhelming majority (93%) of the computed \hat{S}_{22} were found to equal 2 which corresponds to an MA(1) structure. Figure 2 (b) shows a plot of \hat{S}_{22} as computed over the Monte Carlo iterations that more clearly shows the bandwidth estimation procedure in action.

Note that the entries of Table 1a corresponding to the matrix \hat{F}^+ are nearly identical to those of \hat{F} indicating an extremely *low* proportion of \hat{F} matrices that were *not* positive semi-definite even for $T = 100$. The reason for this is twofold: (i) the target values (of those eigenvalues) are relatively large, i.e., not close to zero, and (ii) the bandwidths chosen were appropriate resulting in accurate estimators.

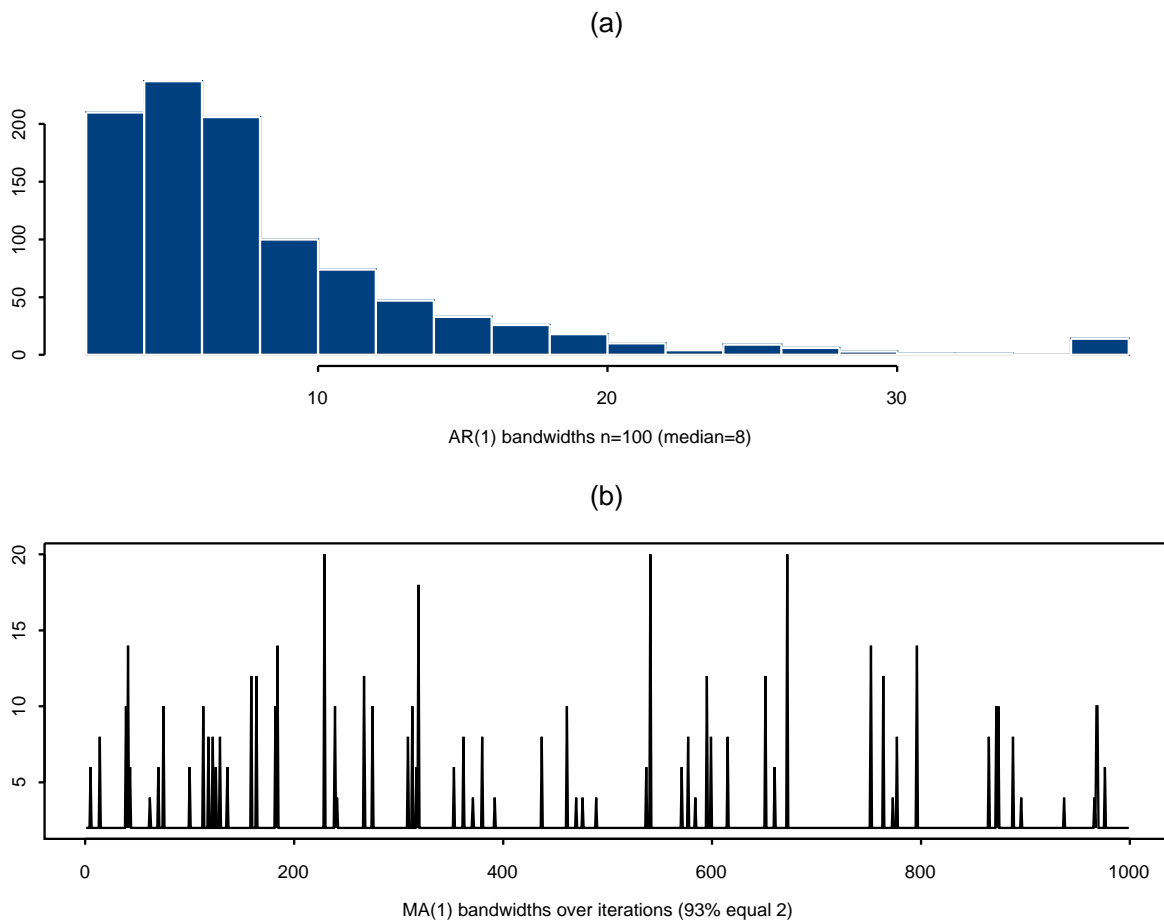


Figure 2: (a) Histogram of \hat{S}_{11} ; (b) Plot of \hat{S}_{22} over the Monte Carlo iterations; Model I with $T=100$.

	\hat{F}_{11}	\hat{F}_{12}	\hat{F}_{22}	\hat{F}_{11}^+	\hat{F}_{12}^+	\hat{F}_{22}^+
κ_B (Bartlett)	1.09	0.62	0.76	n/a	n/a	n/a
κ_{PR} (Parzen)	1.03	1.06	1.05	n/a	n/a	n/a
κ_{trunc} (Truncated-A)	0.99	0.83	0.95	n/a	n/a	n/a
κ_{trunc} (Truncated-E)	1.19	0.42	0.43	1.19	0.42	0.43
$\lambda_{TR,1/2}$ (Trapezoid)	1.22	0.54	0.50	1.22	0.53	0.50
$\lambda_{PR,3/4}$ (Flat-top Parzen)	1.21	0.64	0.65	1.21	0.64	0.65
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	1.23	0.58	0.53	1.22	0.57	0.53
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	1.19	0.69	0.66	1.19	0.68	0.66

Table 1a. Entries represent the empirical MSEs of different estimators relative to (i.e., divided by) the MSE of the optimal second order estimator with kernel κ_{QS} ; Model I with $T = 100$. [Minimum MSE is indicated by boldface.]

	\hat{F}_{11} or \hat{F}_{11}^+	\hat{F}_{12} or \hat{F}_{12}^+	\hat{F}_{22} or \hat{F}_{22}^+
κ_B (Bartlett)	1.24	0.79	0.85
κ_{PR} (Parzen)	1.02	1.11	1.08
κ_{trunc} (Truncated-A)	0.96	0.84	0.94
κ_{trunc} (Truncated-E)	1.13	0.25	0.24
$\lambda_{TR,1/2}$ (Trapezoid)	1.14	0.33	0.27
$\lambda_{PR,3/4}$ (Flat-top Parzen)	1.14	0.33	0.27
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	1.16	0.36	0.29
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	1.19	0.48	0.37

Table 1b. Entries represent the empirical MSEs of different estimators relative to the MSE of the optimal second order estimator with kernel κ_{QS} ; Model I with $T = 500$.

Table 1b is the same as Table 1a with the sample size increased to 500, and the results are qualitatively similar. Notable is the dramatic reduction of all flat-top MSEs of F_{12} and F_{22} going from Table 1a to Table 1b. Interestingly, in Table 1b the MSEs of the flat-top \hat{F}^+ were found *identical* (to 8 decimal points) to those of \hat{F} indicating that there were absolutely *no* occurrences of estimators with negative eigenvalues with the increased sample size; this empirical finding gives credence to the discussion at the end of Section 4.

In order to really see the effect/improvement of \hat{F}^+ vs. \hat{F} , we need to consider a model where the target eigenvalues happen to be close to zero. Model II is characterized by negative (i.e., alternating) dependence which has as its consequence small values for the spectral density at the origin. As a matter of fact, $F_{11}(0)$ is identically zero whereas $F_{22}(0)$ equals 0.052. Coordinate $V_t^{(1)}$ follows an MA(1) model, and $V_t^{(2)}$ follows an ARMA(1,2) model that is—by construction—dependent to coordinate $V_t^{(1)}$ as their cross-correlation is significant at lags around 7. For this reason $F_{12}(w)$ is not identically zero, and the optimal values for S_{12} are not trivial as in Model I; note, however, that $F_{12}(0) = 0$ which is not surprising in view of $F_{11}(0) = 0$ and the Cauchy-Schwarz inequality.

Figure 3 (a) shows histograms of the distribution of the Bartlett estimator of F_{11} in the case of Model II with $T=250$; Figure 3 (b) is the same but concerning the trapezoidal $\lambda_{TR,1/2}$ estimator. As expected, the positivity of the Bartlett estimator results into significant bias when the target value is zero. By contrast, the trapezoidal shows minimal bias albeit somewhat larger variance. But even the variance discrepancy is corrected after the positive-part of the trapezoidal estimator is taken strongly suggesting that the flat-top \hat{F}^+ is a superior estimator. A similar phenomenon occurs with a target value near zero as in the estimation of F_{22} that equals 0.052; see Figure 4.

Table 2a shows the empirically found Mean Squared Errors (MSE) of different estimators relative to the MSE of the kernel κ_{QS} with data from Model II with $T = 100$. The first striking feature of Table 2a is that, despite its optimality among second order kernels, kernel κ_{QS} is vastly outperformed by the traditional positive kernels: Bartlett and Parzen. Those in turn are outperformed by any of our four recommended flat-top kernels in their positive semi-definite variation \hat{F}^+ . The (non-recommended) truncated kernel performs rather erratically regardless of bandwidth choice.

As mentioned above, Model II presents a bit of a challenge in estimating S_{12} by our empirical rule and this difficulty is manifested in the results of Table 2a. The reason for this is that whereas $F_{11}(w)$ equals a constant plus a cosine of period 2π , $F_{12}(w)$ involves a cosine of period $2\pi/7$, i.e., it is very ‘wiggly’. Still, the best flat-top performers, the flat-top Parzen and the infinitely differentiable, manage to achieve a MSE that is about a half of that of the reference kernel κ_{QS} . The situation is dramatically improved if the sample size is increased to 500 as Table 2b shows.

Looking at our four flat-top kernels, $\lambda_{TR,1/2}$, $\lambda_{PR,3/4}$, $\lambda_{QS,4,1}$, and $\lambda_{ID,1/4,0.05}$, the improvement offered by the increased sample size of Table 2b is very apparent, and this is in

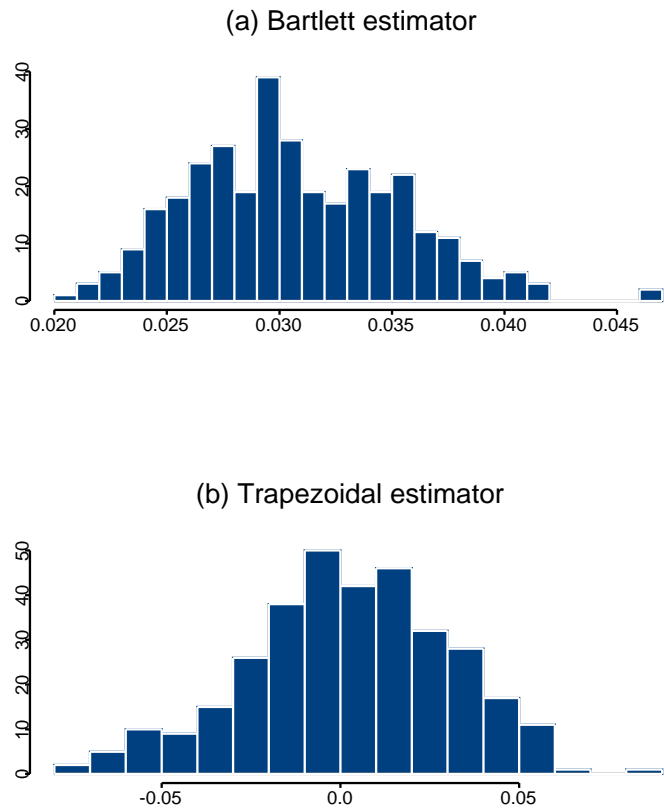


Figure 3: (a) Bartlett estimator of F_{11} ; (b) Trapezoidal estimator of F_{11} ; Model II with $T=250$.




Figure 4: (a) Bartlett estimator of F_{22} ; (b) Trapezoidal estimator of F_{22} ; Model II with $T=250$.

good part due to the bandwidths being chosen by our empirical rule which adapts to the underlying correlation structure. Of course, to realize/maximize those gains, one has to employ the \hat{F}^+ estimators. As conjectured in Section 5, the infinitely differentiable flat-top kernel $\lambda_{ID,1/4,0.05}$ is best overall but with the flat-top Parzen coming in as a (very) close second. Both have impressively low MSEs of the order of 10% as compared to the optimal second order kernel κ_{QS} .

	\hat{F}_{11}	\hat{F}_{12}	\hat{F}_{22}	\hat{F}_{11}^+	\hat{F}_{12}^+	\hat{F}_{22}^+
κ_B (Bartlett)	0.35	0.64	0.61	n/a	n/a	n/a
κ_{PR} (Parzen)	0.81	0.65	0.85	n/a	n/a	n/a
κ_{trunc} (Truncated-A)	0.36	10.5	9.35	n/a	n/a	n/a
κ_{trunc} (Truncated-E)	0.41	19.5	5.10	0.38	5.81	2.25
$\lambda_{TR,1/2}$ (Trapezoid)	0.32	3.06	1.85	0.21	1.18	0.32
$\lambda_{PR,3/4}$ (Flat-top Parzen)	0.10	0.96	0.29	0.07	0.61	0.17
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	0.25	2.37	1.41	0.16	1.00	0.24
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	0.09	1.00	0.28	0.06	0.65	0.16

Table 2a. Entries represent the empirical MSEs of different estimators relative to the MSE of the optimal second order estimator with kernel κ_{QS} ; Model II with $T = 100$.

	\hat{F}_{11}	\hat{F}_{12}	\hat{F}_{22}	\hat{F}_{11}^+	\hat{F}_{12}^+	\hat{F}_{22}^+
κ_B (Bartlett)	0.37	0.09	0.74	n/a	n/a	n/a
κ_{PR} (Parzen)	0.86	0.31	0.95	n/a	n/a	n/a
κ_{trunc} (Truncated-A)	0.24	10.7	30.9	n/a	n/a	n/a
κ_{trunc} (Truncated-E)	0.28	9.30	5.59	0.18	4.25	3.42
$\lambda_{TR,1/2}$ (Trapezoid)	0.23	0.20	0.33	0.12	0.16	0.30
$\lambda_{PR,3/4}$ (Flat-top Parzen)	0.08	0.15	0.12	0.05	0.13	0.12
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	0.18	0.14	0.15	0.09	0.12	0.13
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	0.07	0.10	0.11	0.04	0.09	0.11

Table 2b. Entries represent the empirical MSEs of different estimators relative to the MSE of the optimal second order estimator with kernel κ_{QS} ; Model II with $T = 500$.

Note that the notorious truncated kernel gives poor results in Table 2b (Model II) even with the adaptive bandwidth choice, i.e., the Truncated-E version,¹⁶ whereas it was the best performer in Table 1 (Model I). It is mixed/incoherent results such as these that turned practitioners away from the truncated kernel early on and made people apprehensive regarding infinite-order kernels in general. However, it is the thesis of this paper that those poor results are not associated with the infinite order *per se* but rather with the lack of smoothness of the truncated kernel.

By contrast, all four of our recommended flat-top kernels of Figure 1 beat the traditional kernels in almost all instances of spectral and cross-spectral estimators considered; the single exception is the AR(1) case F_{11} in Model I, the reason being that in that case the traditional estimators enjoy the benefit of an ultra-accurate, model-based, optimal bandwidth choice from a model that happens to be correct. Given the same benefit, flat-top kernels would do similarly well as the example of Truncated-A in Tables 1a and 1b clearly shows.

Finally, to investigate the performance of the flat-top kernels—and the associated bandwidth choice rule—in a set-up with infinite fourth moments, the following model was considered:

MODEL III: $V_t^{(1)} = V_{t-1}^{(1)} - 0.55V_{t-2}^{(1)} + Z_t/\sqrt{\text{Var}(Z_1)}$ where $Z_t = \sigma_t W_t^{(1)}$ with $\sigma_t^2 = 17 \cdot 10^{-7} + 0.04Z_{t-1}^2 + 0.95\sigma_{t-1}^2$, and $V_t^{(2)} = \sum_{j=0}^{\infty} \psi_j W_{t-j}^{(2)}$ where $\psi_j = (j+1)^{-p}$ with $p = 3$. The error series $\{W_t^{(1)}\}$ is i.i.d. with a Student's t_4 distribution, and independent of $\{W_t^{(2)}\}$ that is i.i.d. with a t_3 distribution; the two t distributions were normalized to variance one.

In other words, $V_t^{(1)}$ is an AR(2) model whose innovations Z_t follow a GARCH(1,1) model. The values of the AR(2) parameters were inspired by an example of Phillips, Sun, and Jin (2006) to give an interesting spectral shape; in particular, $F_{11}(w)$ achieves its maximum at $w = \pm\pi/4$ while $F_{11}(0) = 0.526$. The values for the GARCH parameters were chosen to roughly reflect the values obtained by fitting a GARCH(1,1) model to daily returns of the S&P500 index for the 12-year period 1979—1991.

As apparent, $V_t^{(2)}$ is a linear process with strong dependence (and heavy-tailed errors) chosen to illustrate the polynomial rate of part (i) of Theorem 3.1. To elaborate, with

¹⁶The poor performance of the truncated kernel in an MA(1) case with negative dependence was pointed out by West (1997) who instead proposed a model-based covariance estimator; note that this poor performance is clearly not shared by our recommended flat-top kernels as evidenced by Tables 2a and 2b.

coefficients $\psi_j = (j + 1)^{-p}$, the implication is that $\Gamma_{22}(m) = O(m^{-2p+1})$; hence, we need $p > 1$ for F_{22} to be well-defined. The chosen value of $p = 3$ for the simulation yields $F_{22}(0) = (2\pi)^{-1}[\sum_{j=0}^{\infty} \psi_j]^2 \approx 0.23$, and a corresponding polynomial decay of $\Gamma_{22}(m)$ that falls under the premises of part (i) of Theorem 3.1 with a value of r that is slightly less than 4. Finally, note that for practical reasons in the simulation, the infinite sum in the definition of $V_t^{(2)}$ was approximated by the sum of the first 50 terms.

	\hat{F}_{11}	\hat{F}_{12}	\hat{F}_{22}	\hat{F}_{11}^+	\hat{F}_{12}^+	\hat{F}_{22}^+
κ_B (Bartlett)	1.11	0.73	0.89	n/a	n/a	n/a
κ_{PR} (Parzen)	1.03	1.03	1.00	n/a	n/a	n/a
κ_{trunc} (Truncated-A)	1.48	1.17	1.00	n/a	n/a	n/a
κ_{trunc} (Truncated-E)	1.81	0.54	0.86	1.80	0.54	0.86
$\lambda_{TR,1/2}$ (Trapezoid)	1.79	0.62	0.87	1.79	0.59	0.87
$\lambda_{PR,3/4}$ (Flat-top Parzen)	2.23	0.80	0.94	2.23	0.78	0.94
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	1.85	0.63	0.87	1.84	0.61	0.87
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	2.26	0.80	0.95	2.25	0.79	0.94

Table 3a. Entries represent the empirical MSEs of different estimators relative to the MSE of the optimal second order estimator with kernel κ_{QS} ; Model III with $T = 100$.

	\hat{F}_{11} or \hat{F}_{11}^+	\hat{F}_{12} or \hat{F}_{12}^+	\hat{F}_{22} or \hat{F}_{22}^+
κ_B (Bartlett)	1.05	0.86	0.92
κ_{PR} (Parzen)	0.99	1.05	1.00
κ_{trunc} (Truncated-A)	0.49	1.05	0.98
κ_{trunc} (Truncated-E)	0.87	0.36	0.82
$\lambda_{TR,1/2}$ (Trapezoid)	0.87	0.39	0.82
$\lambda_{PR,3/4}$ (Flat-top Parzen)	0.87	0.54	0.88
$\lambda_{QS,4,1}$ (Flat-top Quadratic)	0.92	0.42	0.83
$\lambda_{ID,1/4,0.05}$ (Flat-top Inf. Diff.)	0.97	0.56	0.88

Table 3b. Entries represent the empirical MSEs of different estimators relative to the MSE of the optimal second order estimator with kernel κ_{QS} ; Model III with $T = 500$.

Tables 3a and 3b present the empirically found relative MSEs from a simulation focusing on Model III. Although the MSE of \hat{F}_{jk} is not well-defined in the absence of finite fourth moments, the *relative* MSE (i.e., ratio of two MSEs obtained by different kernels) is well-defined as it is tantamount to the relative MSE of the two normalized spectral densities (i.e., the ratio of the MSEs of two different estimators of \hat{f}_{jk}). Also note that the occurrences of a non-positive semi-definite flat-top estimated matrix under Model III were of the order of 1% when $T = 100$, but dropping to 0.01 of 1% when $T = 500$ lending support to the discussion of Remark 4.1.

Focusing on F_{11} , a first observation is that the flat-top kernels that use the empirical bandwidth choice rule perform roughly comparably to each other; they are uniformly worse than the finite-order kernels in the $T = 100$ case but become uniformly *better* when the sample size increases to 500. Interestingly, the case of F_{11} with $T = 100$ is the *only* instance in the simulations of this paper where the optimal second order kernel κ_{QS} actually performs optimally!

However, the big news here is the remarkable performance of Truncated-A that is by far the best performer when $T = 500$. Comparing this to Truncated-E, the implication is that the Andrews (1991, footnote 5) heuristically motivated plug-in bandwidth for the truncated kernel performs very well in this case, and really makes the difference. To delve into this phenomenon, note that the plug-in bandwidth has a mean of 5.7 and standard deviation of 0.4 over the simulation. By contrast, the empirical bandwidth choice rule of Section 6 has a mean of 6.0 and standard deviation of about 3.7. Thus, it seems that the heavy tails of $V_t^{(1)}$ give our empirical rule a “run for its money” in severely inflating its variability; in fact, the histogram (not shown) of bandwidths chosen by the empirical rule looks heavy-tailed itself.

While it is intuitive that a slight model mis-specification, e.g. using an AR(1) bandwidth on an AR(2) model, may be preferable to feigning total ignorance as our empirical rule does, it is remarkable that a plug-in bandwidth would fare well in the context of heavy tails where the MSE of \hat{F} is not well-defined. This unexpected good performance of the plug-in bandwidth under a lack of finite fourth moments deserves further discussion that is taken up in Appendix B.

Moving on to the polynomial autocovariance decay case of F_{22} , the Bartlett kernel seems to fare quite well here; in fact, it is the best in the $T = 100$ case. This is hardly surprising since in the case of a slow decay of the autocovariance (and associated lack of high-order

smoothness of the spectral density) the gain from using higher order kernels is expected to be small. What is surprising, though, is that with T as small as 500 the flat-top family is the hands-down winner with the truncated and trapezoid sharing first place.

As mentioned before, the polynomial decay is a worst case scenario in terms of showcasing the benefits of using higher-order kernels. Thus, the success of flat-top kernels in Table 3b can partly be attributed to a success of the empirical bandwidth choice rule in this case; this is made manifest by comparing the performance of Truncated-A to that of Truncated-E in Tables 3a and 3b. It appears that the empirical rule automatically captures the appropriate polynomial rate, and gives the resulting improvement in estimation accuracy.

Finally, the results on estimating F_{12} are another success story for the flat-top kernels and the associated empirical bandwidth choice rule. In this case, $\Gamma_{12}(m) = 0$ for all m since the two coordinates of V_t are independent. The adaptive bandwidth choice rule correctly identifies this cross-covariance structure that fits under the premises of part (iii) of Theorem 3.1 with $q = 0$; hence it is not surprising that the truncated kernel is the best here and that the gains over the best second order kernel are dramatic.

Comparing the performance of Truncated-A to that Truncated-E, it is apparent that the Andrews (1991) plug-in bandwidth choice underperforms here. In view of the good performance of this plug-in bandwidth in estimating F_{11} and its poor performance in estimating F_{22} , the poor performance for F_{12} is an elucidation of the fact that—as discussed in the Introduction—the search for a single, optimal bandwidth for the whole matrix is a compromise. Each element of the target matrix is associated with its own optimal bandwidth, and the gains of using that coordinate-specific bandwidth can be appreciable.

With regards to comparing different coordinate-specific bandwidth rules, note that the Andrews (1991) plug-in bandwidth for the truncated kernel essentially amounts to limiting its rate of convergence to that of a second order kernel, and using a rough approximation to the bias constant $c_{\lambda,w}$ of eq. (33) in Appendix B. Thus, in a heavy-tailed context, the practitioner must make a choice between two evils: Andrews’ imposed limitation on the rate of convergence, and the high variance of the empirical bandwidth rule of Section 6. Our simulation suggests that the former may be preferable in a heavy-tailed setting provided, of course, that the AR approximation to the bias/variance constants is reasonable. If an $\text{AR}(p)$ model is used for that approximation, it is necessary that the underlying true spectral density *can* be reasonably approximated by that of an AR model with order given by the

chosen p . This is the case with the first coordinate of Model III, but it is clearly not the case for the second coordinate. All in all, our empirical bandwidth choice rule of Section 6 seems to be a good, totally automatic and generally applicable bandwidth selection tool.

In conclusion, the asymptotic results of Sections 2–6 on the optimal performance of flat-top kernels have been largely substantiated by the finite-sample simulation of the present section. It is of particular importance that the high accuracy of flat-top kernels (after the proposed positive semi-definite transformation) seems to kick in even in sample sizes as small as $T = 100$ making them a valuable tool for practical use. Finally, note that flat-top kernels have recently been shown to be applicable and well-performing in the estimation of polyspectra such as the bispectrum; see Berg and Politis (2008) for more details.

7.2 Application to hypothesis testing

We now re-visit the testing set-up of Remark 3.2. To fix ideas, suppose we wish to test the null hypothesis $H_0 : a'\mu = 0$ against the alternative $H_1 : a'\mu > 0$ at level $\alpha = 0.05$; here $EX_t = \mu = [\mu^{(1)}, \mu^{(2)}]'$, and a is some 2-dimensional vector of interest. For example, $a' = [1, 0]$ is tantamount to a test for $\mu^{(1)}$, whereas $a' = [1, 1]$ yields a test for $\mu^{(1)} + \mu^{(2)}$.

Assuming $\Omega = 2\pi F(0)$ is positive definite, under usual regularity conditions the Central Limit Theorem (CLT):

$$\sqrt{T} \frac{a'(\bar{X}_T - \mu)}{\sqrt{a'\Omega a}} \xrightarrow{\mathcal{L}} N(0, 1) \quad (29)$$

holds. Therefore, our large-sample, 0.05 level test would reject H_0 whenever

$$a'\bar{X}_T > 1.645 \sqrt{a'\hat{\Omega}a/T} \quad (30)$$

where $\hat{\Omega} = 2\pi\hat{F}(0)$ or $2\pi\hat{F}^+(0)$ when the latter is required. Interestingly, note that the strictly positive definite estimator \hat{F}^ϵ of eq. (23) is *not* required in order to carry out this test since it is not a problem to have the RHS of (30) equal zero; this is good in the sense that the practitioner is not required to choose a value for ϵ_T .

Tables 4-6 give the empirical rejection percentages of the above test obtained by simulation using different kernels to estimate $F(0)$. Models I, II and III were considered with a sample size of $T = 250$, and number of replications $B = 999$ as before. The cases $[1,0]$, $[0,1]$, and $[1,1]$ were considered for a' ; the last case was chosen so that the estimated cross-spectral density plays a role as well. For simplicity, μ was chosen such that $\mu^{(1)} = \mu^{(2)}$ using five different values for the latter.

The entries of Table 4 indicate that under Model I all kernels perform in a roughly comparable way in achieving good power while having a size close to the nominal; in general, the kernels that appear more powerful are the ones whose actual size overshoots the nominal the most, thus making a direct comparison difficult. If one focuses on attaining the nominal size as closely as possible, the clear winner is the infinitely differentiable flat-top kernel $\lambda_{ID,1/4,.05}$ while the worst is the Bartlett kernel κ_B advocated by Newey and West (1987).

$a' =$	[1, 0]					[0, 1]					[1, 1]				
$\mu^{(j)} =$	0	.05	.10	.15	.20	0	.05	.10	.15	.20	0	.05	.10	.15	.20
κ_{QS}	8.1	11.4	15.0	20.5	26.9	5.1	8.4	11.3	15.7	20.3	5.5	9.6	16.9	25.2	36.8
κ_B	9.0	13.0	17.0	22.8	28.4	5.3	8.3	11.5	16.3	20.8	6.2	10.8	17.9	26.8	39.0
κ_{PR}	8.1	11.1	15.1	20.6	26.7	5.2	8.4	11.7	15.8	20.6	5.8	10.2	16.7	25.3	36.7
$\kappa_{trunc-A}$	8.1	11.2	15.6	20.7	26.3	5.1	7.9	11.5	15.6	20.5	5.6	9.4	16.4	25.3	36.6
$\kappa_{trunc-E}$	8.2	11.3	14.9	19.9	26.3	4.5	6.4	9.8	14.5	19.4	5.1	9.1	15.3	24.2	35.5
$\lambda_{TR,1/2}$	7.7	9.8	14.4	18.6	24.4	4.5	6.4	9.8	14.5	19.5	5.0	9.0	14.6	23.0	34.2
$\lambda_{PR,3/4}$	7.4	9.8	14.4	18.3	24.1	4.5	6.8	10.2	14.7	19.6	4.9	9.2	14.5	23.2	35.0
$\lambda_{QS,4,1}$	7.5	9.8	14.4	18.1	24.1	4.5	6.4	9.8	14.6	19.5	4.8	9.1	14.5	22.8	33.9
$\lambda_{ID,1/4,.05}$	7.0	9.4	13.4	18.0	23.9	4.6	6.8	10.2	14.6	19.6	5.0	9.2	14.6	22.9	33.7

Table 4. Entries represent the empirical rejection percentages of the one-sided, 0.05 level test of the null hypothesis $H_0 : a'\mu = 0$ using different kernels; Model I with $T = 250$.

The abysmal size performance of all kernels manifested by the entries of Table 5 in the case $a' = [1, 0]$ under Model II is not a fluke. Since $F_{11}(0) = 0$, the CLT (29) does *not* hold, and consequently our test based on (30) is not an appropriate α -level test. To intuitively explain this phenomenon, recall that, since $F_{11}(0) = 0$, the flat-top kernels will yield a negative $\hat{F}_{11}(0)$ with appreciable frequency (that actually tends to 50% as $T \rightarrow \infty$); in these cases, $\hat{F}_{11}^+(0) = 0$, and the test with critical region given by (30) would reject H_0 whenever $\bar{X}_T > 0$. Thus, it is expected here that the size of the test based on flat-top kernels would tend to 1/2 as $T \rightarrow \infty$. By contrast, the size of the test based on finite-order kernels would tend to zero as $T \rightarrow \infty$, and this convergence seems to manifest itself even with a sample size of $T = 250$.

The ill-conditioned case $a' = [1, 0]$ in Model II was included in order to illustrate the failure of test (30) when the CLT (29) fails, since practitioners use this test as a default.

Another reason for inclusion of this case is to compare it with the case $a' = [0, 1]$ where the spectral density is not zero but very close to zero, and the similar case $a' = [1, 1]$. In the latter two cases, the test is well-defined, and—as Table 5 shows—the flat-top kernels have a definite edge overall. In contrast to the finite-order kernels, the flat-top kernels seem to have the accuracy/sensitivity to distinguish a spectral density that is exactly zero from one that is just close to zero. The flat-top $\lambda_{ID,1/4,.05}$ is again the clear winner with $\lambda_{PR,3/4}$ coming in a close second; these two kernels not only manage to get the size approximately correct, they are extremely powerful as well. Not surprisingly, $\lambda_{ID,1/4,.05}$ and $\lambda_{PR,3/4}$ gave the best performing estimators of F under Model II.

$a' =$	[1, 0]					[0, 1]					[1, 1]				
$\mu^{(j)} =$	0	.02	.04	.06	.08	0	.02	.04	.06	.08	0	.02	.04	.06	.08
κ_{QS}	0.0	0.0	0.2	58.9	99.7	0.8	3.4	9.9	20.8	41.9	0.2	3.0	19.6	60.0	91.3
κ_B	0.0	0.0	20.5	97.6	100	1.2	4.5	12.5	24.8	46.9	0.9	7.4	33.3	74.2	97.1
κ_{PR}	0.0	0.0	0.6	72.0	99.8	0.8	3.4	10.0	21.3	42.0	0.2	4.6	24.5	66.3	94.2
$\kappa_{trunc-A}$	21.8	50.9	73.4	94.6	99.8	43.2	64.4	82.7	93.7	98.7	27.4	63.5	91.1	99.4	100
$\kappa_{trunc-E}$	12.9	47.2	74.3	93.3	99.9	19.9	31.6	42.4	51.7	59.4	12.6	32.0	53.6	75.5	93.3
$\lambda_{TR,1/2}$	18.5	59.3	79.5	96.4	100	13.3	24.6	39.0	58.0	73.8	9.3	32.9	67.9	92.1	99.3
$\lambda_{PR,3/4}$	14.3	60.7	91.7	99.8	100	6.2	14.6	29.3	51.7	71.6	5.8	25.9	67.5	92.4	99.8
$\lambda_{QS,4,1}$	17.3	60.3	82.2	97.9	100	10.4	19.0	32.7	54.6	72.8	7.4	27.7	66.2	92.6	99.6
$\lambda_{ID,1/4,.05}$	15.7	61.4	93.5	100	100	5.5	13.4	27.3	50.3	68.3	5.1	23.9	64.3	92.4	99.6

Table 5. Entries represent the empirical rejection percentages of the one-sided, 0.05 level test of $H_0 : a'\mu = 0$; Model II with $T = 250$.

Finally, the entries of Table 6 indicate that all kernels perform comparably to each other in the cases where $a' = [0, 1]$ or $[1, 1]$ in Model III. In the case $a' = [1, 0]$ the finite-order kernels have an edge in capturing the correct size; the flat-top kernels have an edge in power but this comes hand-in-hand with a size larger than nominal, and thus the comparison is difficult.

All in all, these limited finite-sample simulations on testing are encouraging as it seems that using a highly-accurate spectral density estimator, i.e., an estimator with optimal rate of convergence, results in reasonably good performance of the test.

However, it should be noted that the problem of testing involves different objectives (namely size and power), and, in principle, there is no reason to expect that the ker-

nel/bandwidth combination that is optimal for estimating the large-sample covariance matrix Ω will also be optimal for testing. Sun, Phillips, and Jin (2008) make this point eloquently, and devise a criterion to be optimized that is a weighted combination of type I and II errors, thus taking into account size and power at the same time; see also Gao and Gijbels (2008), and Sun and Phillips (2008) for related discussions.

$a' =$	[1, 0]					[0, 1]					[1, 1]				
$\mu^{(j)} =$	0	.02	.04	.06	.08	0	.02	.04	.06	.08	0	.02	.04	.06	.08
κ_{QS}	5.5	8.3	12.8	18.1	23.6	5.5	9.7	13.8	23.2	32.8	5.7	10.8	17.7	28.2	42.2
κ_B	4.6	8.0	11.5	16.6	21.9	5.6	9.7	14.8	23.3	33.1	5.4	10.2	16.8	27.3	41.9
κ_{PR}	5.6	8.6	13.0	17.7	23.7	5.7	9.8	14.6	23.3	33.2	5.7	10.9	18.0	28.3	42.5
$\kappa_{trunc-A}$	7.6	11.5	15.8	22.2	28.2	5.4	9.2	13.6	21.8	33.0	6.2	11.9	19.8	30.5	46.2
$\kappa_{trunc-E}$	6.9	10.7	14.5	21.3	28.4	5.8	9.8	15.5	23.2	33.1	5.6	11.6	19.6	30.9	44.3
$\lambda_{TR,1/2}$	8.0	11.1	15.6	20.7	28.4	5.8	9.8	15.5	23.2	33.2	5.9	12.1	20.9	32.0	45.9
$\lambda_{PR,3/4}$	8.0	11.2	16.0	21.0	28.5	5.4	9.0	13.8	22.6	33.1	5.7	12.2	20.2	31.0	45.4
$\lambda_{QS,4,1}$	8.0	11.7	15.3	20.5	28.1	5.7	9.9	15.2	23.1	33.4	5.9	12.2	20.2	31.9	45.1
$\lambda_{ID,1/4,.05}$	7.3	10.6	14.6	19.2	27.0	5.4	8.9	13.8	22.5	33.2	5.9	11.9	19.5	29.8	43.9

Table 6. Entries represent the empirical rejection percentages of the one-sided, 0.05 level test of $H_0 : a'\mu = 0$; Model III with $T = 250$.

8 Appendix A: Large-sample covariance matrix estimation

Consider the general framework of Andrews (1991) or Hansen (1992) in which the problem at hand is estimation of the large-sample covariance matrix Ω of the sample mean of a second-order stationary (and weakly dependent) sequence of mean zero random vectors $V_t = V_t(\theta)$, $t = 1, \dots, T$, where V_t takes values in \mathbb{R}^d , i.e., Ω as defined in eq. (3).

Here θ is an unknown parameter assumed to have a \sqrt{T} -consistent estimator $\hat{\theta}$, yielding the estimated sequence $\hat{V}_t = V_t(\hat{\theta})$. As an immediate example, consider the case of unknown mean discussed in Remark 3.1. As before, we can define the autocovariance estimators

$$\hat{\Gamma}(j) = \frac{1}{T} \sum_{t=1}^{T-j} \hat{V}_t \hat{V}'_{t+j} \text{ for } j \geq 0, \text{ and } \hat{\Gamma}(j) = \hat{\Gamma}(-j)' \text{ for } j < 0.$$

As usual, we set $\hat{\Gamma}(j) = 0$ for $|j| \geq T$.

The typical heteroskedasticity and autocorrelation consistent (HAC) kernel estimator of Ω has the form

$$\hat{\hat{\Omega}} = \sum_{j=-T}^T \kappa(j/s_T) \hat{\hat{\Gamma}}(j),$$

where the kernel $\kappa(\cdot)$ and the bandwidth parameter $s_T \in [1, T]$ satisfy some standard conditions. A typical condition on κ is:

$\{\kappa : \mathbb{R} \rightarrow [-1, 1], \kappa$ is symmetric, continuous at 0 and for all but a finite number of points,

$$\text{and satisfying } \kappa(0) = 1 \text{ and } \int_{\mathbb{R}} \kappa^2(x) dx < \infty\}. \quad (31)$$

The kernel $\kappa(\cdot)$ is called a ‘spectral window generator’ by Andrews (1991) as it corresponds to the function $K(w) = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} \kappa(j) e^{-ijw}$ that is useful for smoothing the periodogram; here $i = \sqrt{-1}$. In statistics, $\kappa(\cdot)$ is typically called a ‘lag-window’. With the exception of the ‘truncated’ window $\kappa_{trunc}(x)$, the kernels considered by Andrews (1991) and Newey and West (1987) are positive semi-definite, i.e., their respective spectral window $K(w)$ is a nonnegative function.

We now consider the idealized estimator

$$\hat{\hat{\Omega}} = \sum_{j=-T}^T \kappa(j/s_T) \hat{\hat{\Gamma}}(j), \quad (32)$$

that is computed as if the sequence $V_t, t = 1, \dots, T$ were directly observable; the definition of $\hat{\hat{\Gamma}}(m)$ for the above is found in eq. (7).

As in the simple example of unknown mean discussed in Remark 3.1, the estimators $\hat{\hat{\Omega}}$ and $\hat{\hat{\hat{\Omega}}}$ are asymptotically equivalent under general conditions such as Assumptions A, B and C of Andrews (1991) or Condition (V2) of Hansen (1992); see e.g. Theorem 1(b) of Andrews (1991). Intuitively, this is due to the slower rate of convergence of both $\hat{\hat{\Omega}}$ and $\hat{\hat{\hat{\Omega}}}$ as compared to the \sqrt{T} -consistency of $\hat{\theta}$ and $V_t(\hat{\theta})$.

In order to be able to use results such as Theorem 2.1 (ii) and (iii) in the setting of large-sample HAC covariance matrix estimation, we now give a slight generalization of Theorem 1(b) of Andrews (1991) to cover a possible choice of the bandwidth parameter s_T that does not necessarily tend to infinity (or it does at a slow, logarithmic rate).

Lemma 8.1 *Assume Assumptions A, B and C of Andrews (1991) hold true, and that κ satisfies eq. (31). Further assume that, as $T \rightarrow \infty$, we have $s_T/T \rightarrow 0$ and that:*

- (i) $s_T^{-1} \sum_{j=-T+1}^{T-1} |\kappa(j/s_T)| = O(1)$;
 - (ii) $\text{Bias}(\hat{\Omega}) = O(\sqrt{s_T/T})$; and
 - (iii) $s_T \rightarrow \infty$ or $EV_t \frac{\partial}{\partial \theta} V_{t-j} = 0$ for all j .
- Then, $\hat{\Omega} = \Omega + O_P(\sqrt{s_T/T})$, $\hat{\Omega} = \Omega + O_P(\sqrt{s_T/T})$, and $\hat{\Omega} - \hat{\Omega} = o_P(\sqrt{s_T/T})$.

Condition (i) of Lemma 8.1 is immediately satisfied if the kernel κ ‘cuts-off’, e.g., if $\kappa(x) = 0$ for $|x| > \text{some } x_0$. Condition (ii) of Lemma 8.1 can be viewed as a restriction (a lower bound) on the rate of growth of s_T .

Note that the flat-top family of kernels (4) satisfies eq. (31). So, let $\hat{\Omega}_\lambda$ be the estimator $\hat{\Omega}$ that uses a flat-top kernel λ instead of κ , and let $\hat{\hat{\Omega}}_\lambda$ denote its corresponding HAC estimator. In view of Lemma 8.1, the HAC flat-top estimator $\hat{\hat{\Omega}}_\lambda$ inherits the large-sample properties of $\hat{\Omega}_\lambda$ derived in the main body of the paper.

9 Appendix B: Robustness of plug-in bandwidth estimators

In Section 7, the surprisingly good performance of the plug-in bandwidth estimators under a lack of finite fourth moments was observed and deserves further study. To simplify the exposition, we now focus on the one dimensional case of $d = 1$, i.e., bandwidth choice for estimating the spectral density $F(w)$ of the real-valued time series $\{V_t\}$ based on data V_1, \dots, V_T . Thus, let $\hat{F}(w) = \frac{1}{2\pi} \sum_{m=-T}^T \kappa(m/S_T) \hat{\Gamma}(m) e^{-imw}$ where $\hat{\Gamma}(m)$ was defined in (7).

Under regularity conditions—including a finite fourth moment for V_t —we have:

$$\text{Var}(\hat{F}(w)) \sim C_{\kappa,w} F^2(w) \frac{S_T}{T}, \quad \text{and} \quad \text{Bias}(\hat{F}(w)) \sim c_{\kappa,w} F^{(r')}(w) S_T^{-r'} \quad (33)$$

where r' is the minimum of r appearing in part (i) of Theorem 3.1 and the order of the kernel κ , i.e., one for Bartlett’s kernel, two for the usual semi-positive definite kernels, and infinity for the flat-top kernels; see e.g. Rosenblatt (1985). In the above, the quantities $C_{\kappa,w}$ and $c_{\kappa,w}$ depend on w and on the shape of the kernel κ but *not* on F ; hence, they are known to the practitioner.

Nevertheless, F is unknown; so let $\bar{F}(w), \bar{F}^{(r')}(w)$ be preliminary—but consistent—estimators of $F(w), F^{(r')}(w)$ respectively. Typical choices for $\bar{F}(w)$ include a nonparametric kernel estimator such as these discussed in this paper, or the spectral density of a fitted AR(p) model as discussed by Andrews (1991). Often, $\bar{F}^{(r')}(\cdot)$ is obtained by differentiating

r' times the function $\bar{F}(\cdot)$; the consistency of such an estimate based on differentiation is easy to show—see e.g. Politis (2003).

Eq. (33) with $\bar{F}(w)$, $\bar{F}^{(r')}(w)$ plugged-in in place of $F(w)$, $F^{(r')}(w)$ implies a large-sample estimator for the MSE of $\hat{F}(w)$ given by

$$MSE(\hat{F}(w)) \sim C_{\kappa,w} \bar{F}^2(w) \frac{S_T}{T} + [c_{\kappa,w} \bar{F}^{(r')}(w) S_T^{-r'}]^2. \quad (34)$$

Finally, the ‘plug-in’ bandwidth estimator¹⁷ \bar{S}_T is the minimizer of the above MSE expression. Simple algebra shows that

$$\bar{S}_T = \left(\frac{2r' c_{\kappa,w}^2 T}{C_{\kappa,w}} \right)^{1/(2r'+1)} \left(\frac{\bar{F}^{(r')}(w)}{\bar{F}(w)} \right)^{2/(2r'+1)} \quad (35)$$

Note that the plug-in bandwidth estimator \bar{S}_T depends on F *only* through the ratio $\bar{F}^{(r')}(w)/\bar{F}(w)$; this observation has the important implication that \bar{S}_T is *scale-invariant*, i.e., \bar{S}_T remains unchanged if the data V_1, \dots, V_T are all multiplied by the same constant. Consequently, one may conjecture that \bar{S}_T is not adversely affected by possible heavy tails of V_t . To a large extent, this seems to be true; to see this, recall the definition of the normalized spectral density $f(w)$ which, in the $d = 1$ case, reduces to $f(w) = F(w)/\Gamma(0)$ that is estimated by $\bar{f}(w) = \bar{F}(w)/\hat{\Gamma}(0)$ with $f^{(r')}(w)$ estimated by $\bar{f}^{(r')}(w) = \bar{F}^{(r')}(w)/\hat{\Gamma}(0)$.

By the cancellation of $\hat{\Gamma}(0)$ from the ratio $\bar{F}^{(r')}(w)/\bar{F}(w)$, eq. (35) implies

$$\bar{S}_T = \left(\frac{2r' c_{\kappa,w}^2 T}{C_{\kappa,w}} \right)^{1/(2r'+1)} \left(\frac{\bar{f}^{(r')}(w)}{\bar{f}(w)} \right)^{2/(2r'+1)} \quad (36)$$

from which it is apparent that \bar{S}_T is well-defined even in the heavy-tailed case, and the following corollary is immediate.

Corollary 9.1 *Assume $EV_t^2 < \infty$, $\hat{\Gamma}(0) \xrightarrow{P} \Gamma(0) > 0$, and that $\sum_{m=-\infty}^{\infty} |m|^r |\Gamma(m)| < \infty$. Also assume $f(w) > 0$, and let $\bar{f}(w)$, $\bar{f}^{(r')}(w)$ be consistent estimators of $f(w)$, $f^{(r')}(w)$ respectively where r' is the minimum of r and the order of the kernel κ . Then, as $T \rightarrow \infty$,*

$$T^{-1/(2r'+1)} \bar{S}_T \xrightarrow{P} \left(\frac{2r' c_{\kappa,w}^2}{C_{\kappa,w}} \right)^{1/(2r'+1)} \left(\frac{f^{(r')}(w)}{f(w)} \right)^{2/(2r'+1)}. \quad (37)$$

¹⁷The estimated bandwidth \bar{S}_T depends on w but this dependence will not be explicitly denoted.

Note that if the finite fourth moment assumption breaks down, \bar{S}_T as given by (35) or (36) is not the minimizer of $MSE(\hat{F}(w))$ since the latter is not well-defined.¹⁸ Nevertheless, the implication of Corollary 9.1 is that \bar{S}_T grows at the exact rate $T^{1/(2r'+1)}$ even under a potential absence of finite fourth moments, i.e., under the context of Theorem 3.1. This property gives \bar{S}_T a certain robustness for validity. In particular, \bar{S}_T remains a reasonable bandwidth choice even in the heavy tailed set-up since it is well-defined and achieves the desired rate.

The potential disadvantages of the plug-in method, however, have to do with the (in)accuracy of the ‘pilot’ estimator $\bar{F}(w)$ that is plugged in eq. (34). If a nonparametric kernel estimate is employed as a pilot, then a bandwidth choice for the pilot is needed; a poor bandwidth choice for the pilot has serious repercussions on the accuracy of (34) as an estimator of the MSE and, consequently, of \bar{S}_T as an estimator of the optimal bandwidth.

To by-pass this difficulty, Andrews (1991) advocated using the spectral density of a fitted AR(p) model as the pilot. This method seems to work well in practice *provided* that the underlying true spectral density can be well approximated by that of an AR model with order given by the chosen p ; the choice of p , however, is a difficult problem reminiscent of the bandwidth choice problem; an information criterion such as the AIC may be useful in that respect but the practitioner should be aware that the AIC’s objective is not the optimization of spectral density estimation *per se*—see Choi (1992) for more details.

In addition, the application of \bar{S}_T in estimating the bandwidth of a *flat-top kernel* encounters two difficulties: (a) r is unknown, and therefore r' is also unknown;¹⁹ (b) the exact computation of the bias constant $c_{\lambda,w}$ is elusive. Of the two problems, the first one is the most serious—and most difficult—, since it dictates the optimal rate of growth of S_T . Problem (b) has to do with the proportionality constant in S_T that becomes of interest only if the rate of growth of S_T has been identified. For this reason, the empirical bandwidth choice rule of Section 6 focuses on problem (a), i.e., identifying the optimal rate of growth.

Finally, note that the plug-in rule heuristically recommended by Andrews (1991, footnote 5) for the truncated kernel—which is a member of the flat-top family—essentially

¹⁸In general, \bar{S}_T is the minimizer of the RHS of (34) which is an approximation to $MSE(\hat{F}(w))$ only when $EV_t^4 < \infty$.

¹⁹Recall that for finite-order kernels, r' is typically equal to the order of the kernel since more often than not it is smaller than the (unknown) degree of smoothness r .

amounts to limiting its rate of convergence to that of a second order kernel, and using a rough approximation to the bias constant $c_{\lambda,w}$.

10 Technical Appendix

PROOF OF THEOREM 2.1. In view of eq. (8), the proof amounts to showing that

$$\text{Bias}(\hat{F}_{jk}) = \begin{cases} O(1/S_{jk}^r) & \text{in part (i)} \\ O(e^{-acS_{jk}}) + O(1/T) & \text{in part (ii)} \\ O(1/T) & \text{in part (iii)}. \end{cases} \quad (38)$$

To show the above, note that $E\hat{\Gamma}_{jk}(m) = (1 - \frac{|m|}{T})\Gamma_{jk}(m)$. Thus, we have

$$\text{Bias}(\hat{F}_{jk}) \equiv E\hat{F}_{jk} - F_{jk} = A_1 + A_2 + A_3$$

where

$$\begin{aligned} A_1 &= \frac{1}{2\pi} \sum_{m=-T+1}^{T-1} \left(\lambda_{g,c}\left(\frac{m}{S_{jk}}\right) - 1 \right) \Gamma_{jk}(m) e^{-imw} \\ A_2 &= -\frac{1}{2\pi T} \sum_{m=-T+1}^{T-1} |m| \lambda_{g,c}\left(\frac{m}{S_{jk}}\right) \Gamma_{jk}(m) e^{-imw} \\ A_3 &= -\frac{1}{2\pi} \sum_{|m| \geq T} \Gamma_{jk}(m) e^{-imw}. \end{aligned}$$

But $|A_3| \leq \frac{1}{2\pi} \sum_{|m| \geq T} |\Gamma_{jk}(m)| \leq \frac{1}{2\pi T} \sum_{|m| \geq T} |m| |\Gamma_{jk}(m)| = o(1/T)$, since under any of the three conditions (i), (ii) or (iii) we have $\sum_m |m| |\Gamma_{jk}(m)| < \infty$.

Similarly, $|A_2| = O(1/T)$, using the fact that $|\lambda_{g,c}(\frac{m}{S_{jk}})| \leq 1$.

Now note that $A_1 = a_1 + a_2$, where

$$\begin{aligned} a_1 &= \frac{1}{2\pi} \sum_{|m| \leq cS_{jk}} \left(\lambda_{g,c}\left(\frac{m}{S_{jk}}\right) - 1 \right) \Gamma_{jk}(m) e^{-imw} \\ a_2 &= \frac{1}{2\pi} \sum_{cS_{jk} < |m| \leq T} \left(\lambda_{g,c}\left(\frac{m}{S_{jk}}\right) - 1 \right) \Gamma_{jk}(m) e^{-imw} \end{aligned}$$

First observe that $a_1 = 0$, because $\lambda_{g,c}(\frac{m}{S_{jk}}) = 1$ for $|m| \leq cS_{jk}$. Now

$$|a_2| \leq \frac{1}{\pi} \sum_{cS_{jk} < m \leq T} \left| \lambda_{g,c}\left(\frac{m}{S_{jk}}\right) - 1 \right| |\Gamma_{jk}(m)| \leq \frac{1}{\pi} \sum_{cS_{jk} < m \leq T} 2|\Gamma_{jk}(m)| \quad (39)$$

But under the condition of part (i), we have:

$$|a_2| \leq \frac{1}{\pi} \sum_{cS_{jk} < m \leq T} 2 \frac{m^r}{c^r S_{jk}^r} |\Gamma_{jk}(m)| \quad \text{i.e.} \quad \text{Bias}(\hat{F}_{jk}) = O(1/S_{jk}^r) + O(1/T) = O(1/S_{jk}^r).$$

Under the condition of part (ii), eq. (39) gives

$$|a_2| \leq \frac{2C}{\pi} \sum_{cS_{jk} < m \leq T} e^{-am},$$

i.e., $\text{Bias}(\hat{F}_{jk}) = O(e^{-acS_{jk}}) + O(1/T) = O(1/T)$.

Finally, under the condition of part (iii), we have $a_2 = 0$, i.e., $\text{Bias}(\hat{F}_{jk}) = O(1/T)$, and the theorem is proven. \square

For the proof of Theorem 3.1, we will need the following auxiliary lemma.

Lemma 10.1 *Eq. (14), together with the assumption $\Gamma_{jj}(0) > 0$ for all j , implies that*

$$E \left| \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} \right|^{1+\delta} = O(1/T^{\alpha(1+\delta)}) \quad \text{for all } j, k. \quad (40)$$

PROOF OF LEMMA 10.1. Let $\Delta = 1 + \delta$, and note that:

$$\begin{aligned} & E \left| \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} \right|^\Delta = \\ &= E \left| \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\hat{\Gamma}_{kk}(0)} + \sqrt{\Gamma_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} \right|^\Delta \\ &= E \left| \sqrt{\hat{\Gamma}_{kk}(0)}(\sqrt{\hat{\Gamma}_{jj}(0)} - \sqrt{\Gamma_{jj}(0)}) + \sqrt{\Gamma_{jj}(0)}(\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)}) \right|^\Delta \leq c_1 A_1 + c_2 A_2 \end{aligned}$$

where c_1, c_2 are some positive constants. In the above, the simple inequality $(a+b)^\Delta \leq 2^\Delta \max(a, b)^\Delta \leq 2^\Delta(a^\Delta + b^\Delta)$ for $a, b \geq 0$ is used, and

$$A_1 = E \sqrt{\hat{\Gamma}_{kk}(0)^\Delta} |\sqrt{\hat{\Gamma}_{jj}(0)} - \sqrt{\Gamma_{jj}(0)}|^\Delta \quad \text{and} \quad A_2 = \sqrt{\Gamma_{jj}(0)^\Delta} E |\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)}|^\Delta.$$

But $\left(\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)} \right)^\Delta \left(\sqrt{\hat{\Gamma}_{kk}(0)} + \sqrt{\Gamma_{kk}(0)} \right)^\Delta = \left(\hat{\Gamma}_{kk}(0) - \Gamma_{kk}(0) \right)^\Delta$, hence

$$E |\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)}|^\Delta = E \frac{|\hat{\Gamma}_{kk}(0) - \Gamma_{kk}(0)|^\Delta}{\left(\sqrt{\hat{\Gamma}_{kk}(0)} + \sqrt{\Gamma_{kk}(0)} \right)^\Delta} \leq E \frac{|\hat{\Gamma}_{kk}(0) - \Gamma_{kk}(0)|^\Delta}{\sqrt{\Gamma_{kk}(0)^\Delta}} = O(1/T^{\alpha\Delta}) \quad (41)$$

by eq. (14). Therefore, $A_2 = O(1/T^{\alpha\Delta})$.

Note that inequality (41) holds for all k ; hence, it follows that

$$A_1 = O(E|\sqrt{\hat{\Gamma}_{jj}(0)} - \sqrt{\Gamma_{jj}(0)}|^\Delta |\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)}|^\Delta) + O(1/T^{\alpha\Delta}).$$

Finally, observe that the function $h(x) = \sqrt{1-x} - (1-\sqrt{x})$ is nonnegative for all $x \in [0, 1]$.

Therefore, for any $a \geq b > 0$, we have: $\sqrt{a} - \sqrt{b} = |\sqrt{a} - \sqrt{b}| \leq \sqrt{a-b} = \sqrt{|a-b|}$.

Using the above, it follows that

$$\begin{aligned} E|\sqrt{\hat{\Gamma}_{jj}(0)} - \sqrt{\Gamma_{jj}(0)}|^\Delta |\sqrt{\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{kk}(0)}|^\Delta &\leq E\sqrt{|\hat{\Gamma}_{jj}(0) - \Gamma_{jj}(0)|^\Delta} \sqrt{|\hat{\Gamma}_{kk}(0) - \Gamma_{kk}(0)|^\Delta} \\ &\leq \sqrt{E|\hat{\Gamma}_{jj}(0) - \Gamma_{jj}(0)|^\Delta E|\hat{\Gamma}_{kk}(0) - \Gamma_{kk}(0)|^\Delta} = O(1/T^{\alpha\Delta}), \end{aligned}$$

the second inequality being the Cauchy-Schwarz, and the last claim due to eq. (14). Hence, $A_1 = O(1/T^{\alpha\Delta})$ as well, and the lemma is proven. \square .

PROOF OF THEOREM 3.1. Let

$$W_T = \hat{F}_{jk} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}\hat{f}_{jk} = \left(\sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} \right) \hat{f}_{jk}.$$

Focusing on integrability of W_T , note that

$$E|W_T|^\Delta \leq \max |\hat{f}_{jk}|^\Delta E \left| \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)} - \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} \right|^\Delta.$$

But

$$|\hat{f}_{jk}| \leq \frac{1}{2\pi} \sum_{m=-T}^T |\lambda_{g,c}(m/S_{jk})| |\hat{\rho}_{jk}(m)| |e^{-imw}| \leq \frac{1}{2\pi} \sum_{m=-T}^T |\lambda_{g,c}(m/S_{jk})| = O(S_{jk}) \quad (42)$$

by assumption (15). Hence, $\max |\hat{f}_{jk}|^\Delta = O(S_{jk}^\Delta)$. Therefore, by eq. (40) we have:

$$E|W_T|^\Delta = O(S_{jk}^\Delta/T^{\alpha\Delta}). \quad (43)$$

Eq. (43) implies that $W_T = O_P(S_{jk}/T^\alpha)$, and hence $\frac{T^\alpha}{S_{jk}}W_T = O_P(1)$ and $\frac{T^\alpha}{S_{jk}\varepsilon_T}W_T = o_P(1)$ where ε_T is any positive sequence that tends to infinity as $T \rightarrow \infty$. Eq. (43) then also implies that $E|W_T|^\Delta = O(S_{jk}^\Delta\varepsilon_T^\Delta/T^{\alpha\Delta})$. Hence the sequence $\frac{T^\alpha}{S_{jk}\varepsilon_T}W_T$ is uniformly integrable, and we have

$$E \frac{T^\alpha}{S_{jk}\varepsilon_T} W_T = o(1) \quad \text{i.e.,} \quad EW_T = o(S_{jk}\varepsilon_T/T^\alpha),$$

and therefore

$$E\hat{F}_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}E\hat{f}_{jk} + o(S_{jk}\varepsilon_T/T^\alpha).$$

However, $F_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}f_{jk}$; hence,

$$Bias(\hat{F}_{jk}) = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} Bias(\hat{f}_{jk}) + o(S_{jk}\varepsilon_T/T^\alpha). \quad (44)$$

But under any of the conditions of parts (i), (ii), or (iii) we have $Bias(\hat{F}_{jk}) \rightarrow 0$ as $T \rightarrow \infty$ from the results of Theorem 2.1. Furthermore, under any of the conditions of parts (i), (ii), or (iii), we have $S_{jk}/T^\alpha = O(T^{-\eta})$ for some $\eta > 0$. Therefore, letting $\varepsilon_T = \log T$, it follows that $Bias(\hat{f}_{jk}) \rightarrow 0$ as $T \rightarrow \infty$. This fact, coupled with eq. (12), implies $\hat{f}_{jk} = f_{jk} + o_P(1)$.

Now (13) follows by eq. (14) using Jensen's and Markov's inequality. By the fact that $\hat{f}_{jk} = O_P(1)$ and (13) it follows that

$$\hat{F}_{jk} = \sqrt{\hat{\Gamma}_{jj}(0)\hat{\Gamma}_{kk}(0)}\hat{f}_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}\hat{f}_{jk} + O_P(1/T^\alpha). \quad (45)$$

and thus $W_T = O_P(1/T^\alpha)$.

Proof of (i) and (ii). By the above, $T^\alpha W_T = O_P(1)$. Since $S_{jk} \rightarrow \infty$, it follows that $\frac{T^\alpha}{S_{jk}}W_T = o_P(1)$. But then eq. (43) implies that the sequence $\frac{T^\alpha}{S_{jk}}W_T$ is uniformly integrable; hence

$$E\frac{T^\alpha}{S_{jk}}W_T = o(1) \quad \text{i.e.,} \quad EW_T = o(S_{jk}/T^\alpha),$$

and therefore

$$E\hat{F}_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}E\hat{f}_{jk} + o(S_{jk}/T^\alpha).$$

However, $F_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)}f_{jk}$; hence,

$$Bias(\hat{F}_{jk}) = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} Bias(\hat{f}_{jk}) + o(S_{jk}/T^\alpha). \quad (46)$$

But from part (i) of Theorem 2.1 we have: $Bias(\hat{F}_{jk}) = O(1/S_{jk}^r)$; it follows that

$$Bias(\hat{f}_{jk}) = O(1/S_{jk}^r) + o(S_{jk}/T^\alpha). \quad (47)$$

Recall that $Var(\hat{f}_{jk}) = O(S_{jk}/T)$ by eq. (12). Note that the second term in $Bias(\hat{f}_{jk})$ is of bigger order than the standard deviation of \hat{f}_{jk} since $\alpha \leq 1/2 \leq (r+1)/(2r+1)$.

Hence, minimization of the order of magnitude of the Mean Squared Error of \hat{f}_{jk} gives the stated optimal choice for the bandwidth S_{jk} in part (i) of Theorem 3.1, and the resulting

rate of convergence of \hat{f}_{jk} as given in eq. (16). Finally, note that the $O_P(1/T^\alpha)$ term in eq. (45) is negligible compared to the accuracy of \hat{f}_{jk} as given in (16). Thus, eq. (45) together with (16) implies (17), and part (i) is proven.

To prove part (ii), recall that from part (ii) of Theorem 2.1 we have $Bias(\hat{F}_{jk}) = O(1/T)$. Plugging the optimal bandwidth $S_{jk} = A \log T$ in eq. (46) we obtain:

$$Bias(\hat{f}_{jk}) = O(1/T) + o(\log T/T^\alpha) = O(\log T/T^\alpha). \quad (48)$$

Recall that $Var(\hat{f}_{jk}) = O(\log T/T)$ by eq. (12). Hence, minimization of the order of magnitude of the Mean Squared Error of \hat{f}_{jk} gives the stated rate of convergence of \hat{f}_{jk} . By eq. (45), \hat{F}_{jk} has the same rate of convergence as \hat{f}_{jk} , and part (ii) is proven.

Proof of (iii). Note that $\frac{T^\alpha}{\varepsilon_T} W_T = o_P(1)$ where ε_T is any positive sequence that tends to infinity as $T \rightarrow \infty$; for example, we could take $\varepsilon_T = \log \log T$. Also note that S_{jk} is constant under the premises of part (iii). Thus, eq. (43) implies $E|T^\alpha W_T|^\Delta = O(1)$, and thus the sequence $\frac{T^\alpha}{\varepsilon_T} W_T$ is uniformly integrable. Hence

$$E \frac{T^\alpha}{\varepsilon_T} W_T = o(1) \quad \text{i.e.,} \quad E W_T = o(\varepsilon_T/T^\alpha),$$

and therefore

$$E \hat{F}_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} E \hat{f}_{jk} + o(\varepsilon_T/T^\alpha).$$

However, $F_{jk} = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} f_{jk}$; hence,

$$Bias(\hat{F}_{jk}) = \sqrt{\Gamma_{jj}(0)\Gamma_{kk}(0)} Bias(\hat{f}_{jk}) + o(\varepsilon_T/T^\alpha).$$

But from part (iii) of Theorem 2.1 we have: $Bias(\hat{F}_{jk}) = O(1/T)$; it follows that

$$Bias(\hat{f}_{jk}) = O(1/T) + o(\varepsilon_T/T^\alpha) = O(\varepsilon_T/T^\alpha). \quad (49)$$

Recalling that $Var(\hat{f}_{jk}) = O(1/T)$ by eq. (12), gives the stated rate of convergence for \hat{f}_{jk} which—by eq. (45)—is the same as that of \hat{F}_{jk} , and part (iii) of the theorem is proven. \square

PROOF OF THEOREM 4.1. The condition $\hat{F} = F + O_P(1/R_T)$ implies

$$\hat{\Lambda} = \Lambda + O_P(1/R_T), \quad \text{and hence} \quad \hat{\lambda}_j = \lambda_j + O_P(1/R_T) \quad \text{for all } j; \quad (50)$$

see e.g. Theorems 3.2 and 4.2 (and the discussion afterwards) of Eaton and Tyler (1991). But, viewed as an estimator of the nonnegative λ_j , $\hat{\lambda}_j^+$ is a better (or, at least, not worse) estimator than $\hat{\lambda}_j$ in the sense that $|\hat{\lambda}_j^+ - \lambda_j| \leq |\hat{\lambda}_j - \lambda_j|$ always. Hence, it follows that

$$\hat{\lambda}_j^+ = \lambda_j + O_P(1/R_T) \text{ for all } j, \text{ and hence } \hat{\Lambda}^+ = \Lambda + O_P(1/R_T). \quad (51)$$

Using eq. (50) and (51) we have the following:

$$\begin{aligned} F + O_P(1/R_T) &= \hat{F} = \hat{U}\hat{\Lambda}\hat{U}^* = \hat{U}(\Lambda + O_P(1/R_T))\hat{U}^* \\ &= \hat{U}(\Lambda^+ + O_P(1/R_T))\hat{U}^* = \hat{F}^+ + O_P(1/R_T), \end{aligned}$$

the latter since $\hat{U} = U + o_P(1) = O_P(1)$; solving for \hat{F}^+ in the above, the theorem is proven. \square

PROOF OF THEOREM 6.1. The proof is analogous to the proof of Theorem 2.3 of Politis (2003) and is omitted. \square

PROOF OF LEMMA 8.1. The case $s_T \rightarrow \infty$ is covered in Theorem 1 of Andrews (1991); thus, we now assume $EV_t \frac{\partial}{\partial \theta} V_{t-j} = 0$ for all j .

A careful reading of the proof of Theorem 1(b) of Andrews (1991) indicates that the proof first hinges on showing that $(Ts_T)^{-1/2} \sum_{j=-T+1}^{T-1} \kappa(|j|/s_T) \rightarrow 0$; but this follows immediately from our condition (i).

Now noting that $T^{-1} \sum_{t=j+1}^T V_t \xrightarrow{P} 0$ from a Weak Law of Large Numbers under Assumption A, we further need to show that $T^{-1} \sum_{t=j+1}^T V_t \frac{\partial}{\partial \theta} V_{t-j} \xrightarrow{P} 0$. But this follows from a Weak Law of Large Numbers for the cross-correlation of the series V_t to the series $\frac{\partial}{\partial \theta} V_{t-j}$ under Assumption C and our assumption $EV_t \frac{\partial}{\partial \theta} V_{t-j} = 0$. \square

References

- [1] Andrews, D. (1991). Heteroskedasticity and autocorrelation consistent covariance matrix estimation, *Econometrica*, **59**, 817-858.
- [2] Andrews, D. (1992). Equivalence of the higher order asymptotic efficiency of k -step and extremum statistics, *Econometric Theory*, **18**, 1040-1085.

- [3] Andrews, D. and Monahan, J. (1992). An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator, *Econometrica*, **60**, 953-966.
- [4] Berg, A. and Politis, D.N. (2008). Higher-order polyspectral estimation with flat-top lag-windows, *Ann. Inst. Statist. Math.* (to appear).
- [5] Brillinger, D.R. (1979). Confidence intervals for the crosscovariance function. *Selecta Statistica Canadiana*, 5, pp. 3-16.
- [6] Brillinger, D.R. (1981), *Time Series: Data Analysis and Theory*, Holden-Day, New York.
- [7] Brillinger, D.R. and Rosenblatt, M. (1967). Asymptotic theory of k th order spectra. In *Spectral Analysis of Time Series*, (B. Harris, Ed.), Wiley, New York, pp. 153-188.
- [8] Brockwell, P. J. and Davis, R. A. (1991), *Time Series: Theory and Methods, 2nd ed.*, Springer, New York.
- [9] Choi, B.-S. (1992). *ARMA Model Identification*, Springer, New York.
- [10] Davis, R. A., and Mikosch, T. (2000). The sample autocorrelations of financial time series models. In *Nonlinear and nonstationary signal processing*, pp. 247-274, Cambridge Univ. Press, Cambridge.
- [11] Donoho, D.L. and Johnstone, I.M. (1995). Adapting to unknown smoothness via wavelet shrinkage, *J. Amer. Statist. Assoc.*, 90, 1200-1224.
- [12] Eaton, M.E. and Tyler, D.E. (1991). On Weilandt's inequality and its application to the asymptotic distribution of the eigenvalues of a random symmetric matrix, *Annals Statist.*, **19**, No. 1, 260-271.
- [13] Embrechts, P., Klüppelberg, C. and Mikosch, T. (1997). *Modelling Extremal Events*. Springer, Berlin.
- [14] Epanechnikov, V.A. (1969). Non-parametric estimation of a multivariate probability density, *Theory of Prob. and its Applications*, vol. 14, 153-158.
- [15] Gallant, A.R. (1987). *Nonlinear Statistical Models*, John Wiley, New York.

- [16] Gallant, A.R. and White, H. (1988). *A Unified Theory of Estimation and Inference for Nonlinear Dynamic Models*, Basil Blackwell, New York.
- [17] Gao, J. and Gijbels, I. (2008). Bandwidth selection in nonparametric kernel testing, *J. Amer. Statist. Assoc.*, 103, 1584-1594.
- [18] Hall, P. and Yao, Q. (2003). Inference in ARCH and GARCH models with heavy-tailed errors, *Econometrica*, 71, 285-317.
- [19] Hamilton, J.D. (1994). *Time Series Analysis*. Princeton University Press, Princeton.
- [20] Hannan, E.J. (1957). The variance of the mean of a stationary process, *J. Roy. Statist. Soc., Ser. B*, 19, 282-285
- [21] Hannan, E.J. (1958). The estimation of the spectral density after trend removal. *J. Roy. Statist. Soc., Ser. B*, 20, 323-333.
- [22] Hannan, E.J. (1970), *Multiple Time Series*, John Wiley, New York.
- [23] Hannan, E.J. and Deistler, M. (1988). *The Statistical Theory of Linear Systems*, John Wiley, New York.
- [24] Hansen, B.E. (1992). Consistent covariance matrix estimation for dependent heterogeneous processes, *Econometrica*, **60**, 967-972.
- [25] Hansen, L.P. (1982). Large sample properties of generalized method of moments estimators, *Econometrica*, **50**, 1029-1054
- [26] Hashimzade, N. and Vogelsang, T.J. (2008). Fixed-b asymptotic approximations of the sampling behavior of nonparametric spectral density estimators. *J. Time Ser. Anal.*, 29, 142–162.
- [27] Inoue, A. and Shintani, M. (2006). Bootstrapping GMM estimators for time series, *J. Econometrics*, 133, 531–555.
- [28] Jowett, G.H. (1954), The comparison of means of sets of observations from sections of independent stochastic series, *J. Roy. Statist. Soc., Ser. B*, 17, 208-227.

- [29] Kiefer, N.M. and Vogelsang, T.J. (2002). Heteroscedasticity-autocorrelation robust standard errors using the Bartlett kernel without truncation, *Econometrica*, 70, 1350-1366.
- [30] Kiefer, N.M. and Vogelsang, T.J. (2005). A new asymptotic theory for heteroscedasticity-autocorrelation robust tests, *Econometric Theory*, 21, 1130-1164.
- [31] Lahiri, S.N. (2003), *Resampling Methods for Dependent Data*, Springer Verlag, New York.
- [32] Leadbetter, M.R., Lindgren, G., and Rootzen, H. (1983), *Extremes and related properties of random sequences and processes*, Springer-Verlag, New York.
- [33] Lepski, O.V. (1990). On a problem of adaptive estimation in Gaussian white noise, *Th. Prob. Appl.*, 35, 454-466.
- [34] Lepski, O.V. (1991). Asymptotically minimax adaptive estimation I: Upper bounds, optimally adaptive estimates, *Th. Prob. Appl.*, 36, 682-697.
- [35] Lepski, O.V. (1992). Asymptotically minimax adaptive estimation II: Schemes without optimal adaptation, *Th. Prob. Appl.*, 37, 433-468.
- [36] McMurry, T. and Politis, D.N. (2004). Nonparametric regression with infinite order flat-top kernels, *J. Nonparam. Statist.*, vol. 16, no. 3-4, 549-562.
- [37] Newey, W. and West, K. (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica*, 55, 703-708.
- [38] Newey, W. and West, K. (1994). Automatic lag selection in covariance matrix estimation, *Rev. Econ. Studies*, 61, 631-653.
- [39] Ng, S. and Perron, P. (1996). The exact error in estimating the spectral density at the origin, *J. Time Ser. Anal.*, vol. 17, no. 4, 379-408.
- [40] Parzen, E. (1957). On Consistent Estimates of the Spectrum of a Stationary Time Series. *Ann. Math. Statist.*, Vol. 28, No. 2., pp. 329-348.
- [41] Parzen, E. (1961), Mathematical Considerations in the Estimation of Spectra, *Technometrics*, vol. 3, 167-190.

- [42] Phillips, P.C.B., Sun, Y. and Jin, S. (2006). Spectral density estimation and robust hypothesis testing using steep origin kernels without truncation. *Intern. Econ. Review*), 47, 837-894.
- [43] Politis, D.N. (2001). On nonparametric function estimation with infinite-order flat-top kernels, in *Probability and Statistical Models with applications*, Ch. Charalambides et al. (Eds.), Chapman and Hall/CRC: Boca Raton, pp. 469-483.
- [44] Politis, D.N. (2003). Adaptive bandwidth choice, *J. Nonparam. Statist.*, vol. 15, no. 4-5, 517-533.
- [45] Politis, D.N. (2004). A heavy-tailed distribution for ARCH residuals with application to volatility prediction, *Annals of Economics and Finance*, vol. 5, pp. 283-298.
- [46] Politis, D.N., and Romano, J.P. (1995), Bias-corrected nonparametric spectral estimation. *J. Time Ser. Anal.*, **16**, 67–103.
- [47] Politis, D.N., Romano, J.P. and Wolf, M. (1999), *Subsampling*, Springer Verlag, New York.
- [48] Politis, D.N., and White, H. (2004). Automatic block-length selection for the dependent bootstrap, *Econometric Reviews*, vol. 23, no. 1, pp. 53-70.
- [49] Priestley, M.B. (1962), Basic considerations in the estimation of spectra, *Technometrics*, vol. 4, 551-564.
- [50] Priestley, M.B. (1981). *Spectral Analysis and Time Series*, Adademic Press, New York.
- [51] Robinson, P. (1991). Automatic frequency domain inference on semiparametric and nonparametric models, *Econometrica*, vol. 59, 1329-1363.
- [52] Robinson, P.M. and Velasco, C. (1997). Autocorrelation-robust inference. In *Handbook of Statistics*, (G.S. Maddala and C.R. Rao, Eds.), vol. 15, pp. 267-298, Amsterdam: North Holland.
- [53] Romano, J.P. and Thombs, L. (1996). Inference for autocorrelations under weak assumption, *J. Amer. Statist. Assoc.*, 91, 590-600.

- [54] Rosenblatt, M. (1984). Asymptotic normality, strong mixing and spectral density estimates, *Annals Prob.*, 12, 1167-1180.
- [55] Rosenblatt, M. (1985), *Stationary Sequences and Random Fields*, Birkhäuser, Boston.
- [56] Samarov, A. (1977). Lower bound for the risk of spectral density estimates, *Probl. Inform. Transm.*, 13, pp. 67-72.
- [57] Shao, X. and Wu, W.B. (2007). Asymptotic spectral theory for nonlinear time series. *Annals Statist.*, Vol. 35, 1773-1801.
- [58] Sun, Y., Phillips, P.C.B., and Jin, S. (2008). Optimal bandwidth selection in heteroscedasticity-autocorrelation robust testing, *Econometrica*, 76, pp. 175-194.
- [59] Sun, Y., and Phillips, P.C.B. (2008). Optimal bandwidth choice for interval estimation in GMM regression. Cowles Foundation Discussion Paper No. 1661, Yale University.
- [60] Velasco, C. and Robinson, P.M. (2001). Edgeworth Expansions for Spectral Density Estimates and Studentized Sample Mean, *Econometric Theory*, Vol. 17, pp. 497-539
- [61] West, K.D. (1997), Another heteroskedasticity— and autocorrelation—consistent covariance matrix estimator, *J. Econometrics*, 76, pp. 171-191.