

Projection inference for high-dimensional covariance matrices with structured shrinkage targets

Fabian Mies

e-mail: mies@stochastik.rwth-aachen.de

Ansgar Steland

e-mail: steland@stochastik.rwth-aachen.de

Abstract: Analyzing large samples of high-dimensional data under dependence is a challenging statistical problem as long time series may have change points, most importantly in the mean and the marginal covariances, for which one needs valid tests. Inference for large covariance matrices is especially difficult due to noise accumulation, resulting in singular estimates and poor power of related tests. The singularity of the sample covariance matrix in high dimensions can be overcome by considering a linear combination with a regular, more structured target matrix. This approach is known as shrinkage, and the target matrix is typically of diagonal form. In this paper, we consider covariance shrinkage towards structured nonparametric estimators of the bandable or Toeplitz type, respectively, aiming at improved estimation accuracy and statistical power of tests even under nonstationarity. We derive feasible Gaussian approximation results for bilinear projections of the shrinkage estimators which are valid under nonstationarity and dependence. These approximations especially enable us to formulate a statistical test for structural breaks in the marginal covariance structure of high-dimensional time series without restrictions on the dimension, and which is robust against nonstationarity of nuisance parameter. We show via simulations that shrinkage helps to increase the power of the proposed tests. Moreover, we suggest a data-driven choice of the shrinkage weights, and assess its performance by means of a Monte Carlo study. The results indicate that the proposed shrinkage estimator is superior for non-Toeplitz covariance structures close to fractional Gaussian noise.

MSC2020 subject classifications: Primary 62G05, 60F15; secondary 62M10.

Keywords and phrases: Change-point, Gaussian approximation, nonstationary time series, bilinear form, shrinkage.

Received October 2022.

1. Introduction

Many modern applications encounter large data sets of very high dimension d , which requires new approaches for modeling as well as statistical inference, since classical methods usually fail. A particularly prominent example is the covariance matrix Σ of a random vector of d variables, which has d^2 unknown

parameters. It is well-known that the sample covariance matrix $\hat{\Sigma}_n$ of sample size n is singular if $n < d$, and consistency of $\hat{\Sigma}_n$ for growing dimension d requires to control the $O(d^2)$ covariances and otherwise generally fails. Specifically, for n iid samples of d -dimensional random vectors with finite fourth moments, consistency in terms of the expected squared Frobenius norm of $\hat{\Sigma}_n - \Sigma_n$ requires that $\text{diag}(\Sigma_n)$ and the variance of the average of the squared variables to be of the order $o(n/d)$, (Ledoit and Wolf, 2004). In terms of eigenvalues and eigenvectors, classic inconsistency results for the largest eigenvalue and the top eigenvector under the regime $d/n \rightarrow c \in (0, 1)$ are due to (Johnstone, 2001; Johnstone and Lu, 2009), even when $\Sigma = I$. Since, however, it is generally believed that many high-dimensional data sets are governed by, say $r \ll d$, strong signals, spiked covariance models with r leading eigenvalues strictly larger than 1, whereas the remaining bulk of the spectrum consists of unit eigenvalues, are quite extensively studied. Here the leading eigenvalues of $\hat{\Sigma}_n$ are biased, thus requiring de-biasing procedures which, however, depend on the loss function, and the eigenvectors are inconsistent as well, see (Donoho, Gavish and Johnstone, 2018) and the references therein. Further potential remedies are to impose additional structural assumptions on Σ , e.g. sparsity (Bickel and Levina, 2008a), Toeplitz shape (Cai, Ren and Zhou, 2013), or assuming the entries further from the diagonal to decay rapidly, also known as a bandable covariance matrix (Bickel and Levina, 2008b). In these frameworks, it is possible to derive consistent estimators of Σ even if $d \gg n$, see Cai, Ren and Zhou (2016) for a survey of suitable estimators and matching minimax bounds.

The structured covariance matrix estimators admit optimal rates of convergence within the specified class of matrices. However, if the true Σ does not satisfy the structural assumptions, the estimators might perform poorly. For instance, if we use an estimator of Toeplitz form, but Σ is not a Toeplitz matrix, then this estimator will be inconsistent. In order to strike a compromise between structured and model-free estimation, one may consider a shrinkage estimator of the form

$$\hat{\Sigma}_n^w = (1 - w) \hat{\Sigma}_n + w \tilde{\Sigma}_n, \quad w \in [0, 1],$$

where $\hat{\Sigma}_n$ is the usual sample covariance matrix, and the shrinkage target $\tilde{\Sigma}_n$ is a structured estimator. The weight w may be interpreted as quantifying the confidence in the structural assumption underlying the estimator $\tilde{\Sigma}_n$. A particularly useful feature of shrinkage estimators is that $\hat{\Sigma}_n^w$ will be regular, as $\tilde{\Sigma}_n$ is usually chosen as a regular matrix. This makes shrinkage estimation especially interesting for situations where the inverse of the covariance matrix is required, e.g. in portfolio optimization. Ledoit and Wolf (2004) suggest to use a multiple of the identity matrix, and Steland (2018) studies more general diagonal matrices. Linear shrinking towards a structured target with respect to the scaled Frobenius loss has been studied in (Ledoit and Wolf, 2003), but the results therein on consistent estimators of the optimal shrinking weights consider classical fixed- d asymptotics only. In the present paper, we go beyond a convex combination of the sample covariance matrix with a single shrinkage target. Instead, we study

convex combinations with a structured covariance estimator with known minimax optimality properties for bandable Σ , and a structured covariance estimator which is minimax-optimal if Σ is of Toeplitz type, by considering (if $w > 0$) the parameterized shrinkage target

$$\tilde{\Sigma}_n = \frac{w_2}{w} \hat{\Sigma}_n^\dagger + \frac{w_3}{w} \hat{\Sigma}_n^\circ,$$

for nonparametric estimators $\hat{\Sigma}_n^\dagger$ and $\hat{\Sigma}_n^\circ$, respectively, addressing both types of matrices, for weights $0 \leq w_2, w_3 \leq 1$ such that $w_1 + w_2 + w_3 = 1$ where $w_1 = (1 - w)$.

Theoretical analysis of these structured covariance estimators typically focuses on their consistency and rate of convergence. Here, we contribute by showing that the estimators proposed when Σ is bandable and Toeplitz, respectively, are still consistent under a general nonstationary nonlinear time series model, whereas the known results assume iid d -dimensional random vectors. Thus, by picking (asymptotically) the corresponding unit vectors as weights, the proposed shrinkage estimator can adapt to, firstly, a Toeplitz structure, i.e. stationarity of the coordinates, or, secondly, a bandable covariance structures, and then use consistent estimators with guaranteed good rates of convergence, or, thirdly, use the purely nonparametric sample covariance estimator. Moreover, to perform statistical inference, distributional approximations are required. Here, we consider inference about the high-dimensional covariance matrix in terms of projections $v^T \Sigma v$ for some vector $v \in \mathbb{R}^d$. Our results provide valid distributional approximations even when the choice of the threshold parameter of the estimators addressing Toeplitz-type or bandable covariances does not guarantee consistency. The vector v may either be determined by the application, or chosen at random, independently from the sample. For example, in risk management, v may represent weights of a portfolio of risky assets, and $v^T \Sigma v$ corresponds to the variance of this portfolio. The choice $v = (1, \dots, 1)^T / d$ corresponds to the equal-weighted portfolio which is more diversified compared to portfolios using weights according to the market capitalization such as exchange-traded funds (ETFs) on market indices. A recent study on its performance is [Swade et al. \(2023\)](#). For other approaches one may estimate the weights from historical data and, by conditioning, assume they are fixed. As another example, the hypothesis $\Sigma = \Sigma_0$ may be studied via random projections: If v is drawn from some continuous distribution, then $\Sigma \neq \Sigma_0$ implies $v^T \Sigma v \neq v^T \Sigma_0 v$ almost surely. It is thus reasonable to employ a test statistic of the form $v^T (S - \Sigma_0) v$ for some estimator S , either structured, unstructured, or shrunken. Further applications (lasso prediction, sparse principal components) are discussed below. The idea to study high-dimensional covariance matrices via bilinear forms has been suggested in [Steland and von Sachs \(2017, 2018\)](#), and applied to changepoint testing ([Steland, 2020, 2023](#)), K-sample problems ([Mause and Steland, 2020](#)) and factor models, ([Bours and Steland, 2021](#)). For extensions to methods in the frequency domain in terms of bilinear forms of spectral density matrix estimators and related statistics including nonlinear spectral averages see ([Steland, 2024](#)). In the latter references, as well as in the present paper, the important assumption is

that the projection vectors v have a bounded l_1 -norm. This condition is satisfied in many applications and enables the methodology to work also for very high dimensions d .

In this paper, we derive Gaussian approximation results for shrinkage estimators of the covariance matrix with structured shrinkage targets $\tilde{\Sigma}_n$. Our results are valid for nonstationary, nonlinear, high-dimensional time series X_t , and we demonstrate how to perform inference based on a suitable bootstrap scheme. Notably, we only impose assumptions on the marginal time series, while the dependency between the components may be arbitrary, and we need no restrictions on the dimension d . This is a beneficial consequence of the projection technique, as described in Section 2. As we allow the dimension to grow with sample size, the considered models are arrays of time series. Thus, classical central limit theory is not applicable because a limiting distribution in general does not exist. Instead, we formulate our results in terms of so-called strong approximations, or sequential couplings, in the sense of Komlós, Major and Tusnády (1975). Compared to the existing work of Steland and von Sachs (2017, 2018), we consider a much broader class of nonlinear time series models, which also allows us to study non-diagonal, structured shrinkage targets. Underlying our mathematical results are recent Gaussian couplings for nonstationary time series established in Mies and Steland (2023) for the regime $d \ll n^{\frac{1}{3}}$. We extend these results to projections $v^T X_t$ of very high-dimensional time series X_t . In particular, it is shown that for l_1 -bounded projection vectors v , the Gaussian approximation is valid irrespective of the ambient dimension d , and of the dependence structure between the d components. It turns out that this general projection framework is also applicable to bilinear forms of sample covariance matrices, as demonstrated in Section 3. The presented distributional approximation also holds sequentially, and we demonstrate how to apply this to a test for structural breaks in the covariance matrix. An attractive feature of the proposed test is that it is robust against general nonstationarity under the null hypothesis, and only detects changes in the target parameter, i.e. the marginal covariance structure. Furthermore, we discuss the optimal choice of shrinkage weights and provide a data-driven criterion. Simulation results demonstrate that the asymptotic theory is also applicable in finite samples, and that shrinkage may improve the performance both for estimation and for change testing.

The rest of this paper is structured as follows. In Section 2, we present a general sequential Gaussian approximation result for projections of high-dimensional nonstationary time series, which is applied to bilinear forms of covariance matrices in Section 3. The data-driven choice of shrinkage weights is discussed in Section 4. Section 5 describes the application of our distributional approximations to tests for changes in the covariance structure of a high-dimensional time series, and presents a feasible bootstrap scheme. Simulation results are presented in Section 6.

2. Gaussian approximations for projections

We consider a d -variate time series X_t which may be nonlinear and nonstationary, given by

$$X_t = G_t(\epsilon_t, \epsilon_{t-1}, \dots), \tag{1}$$

for measurable mappings $G_t : \mathbb{R}^\infty \rightarrow \mathbb{R}^d$, $t = 1, \dots, n$, where we endow \mathbb{R}^∞ with the σ -algebra generated by all finite projections. The ϵ_i are iid $U[0, 1]$ random variables, which may be regarded as random seeds for the time series. Note that the distributional assumption on the ϵ_i is not restrictive, as G_t is not necessarily be smooth, but just measurable. The model formulation (1) allows for a convenient formulation of ergodicity conditions via the physical dependence measure introduced by Wu (2005). To formulate our assumptions on the kernels G_t , introduce a second sequence $\tilde{\epsilon}_i$ of iid $U[0, 1]$ random variables, independent of the sequence ϵ_i . For any $t \in \mathbb{N}$, denote

$$\begin{aligned} \epsilon_t &= (\epsilon_t, \epsilon_{t-1}, \dots) \in \mathbb{R}^\infty, \\ \tilde{\epsilon}_{t,j} &= (\epsilon_t, \dots, \epsilon_{j+1}, \tilde{\epsilon}_j, \epsilon_{j-1}, \dots) \in \mathbb{R}^\infty, \end{aligned}$$

such that $X_t = G_t(\epsilon_t) = (G_{t,l}(\epsilon_t))_{l=1}^d$. We assume that the impact of past random seeds decays to zero polynomially, such that for some $\beta > 1$

$$\begin{aligned} (\mathbb{E}|G_{t,l}(\epsilon_t) - G_{t,l}(\tilde{\epsilon}_{t,t-j})|^q)^{\frac{1}{q}} &\leq \tilde{\Theta} j^{-\beta}, \quad j \geq 1, \quad l = 1, \dots, d, \\ (\mathbb{E}|G_{t,l}(\epsilon_0)|^q)^{\frac{1}{q}} &\leq \tilde{\Theta}, \quad l = 1, \dots, d. \end{aligned} \tag{P.1}$$

The time series X_t may be nonstationary, and the nonstationarity is explicit by making the kernel G_t depend on t . Nevertheless, to obtain stronger asymptotic results, we require some regularity in time. Here, we do not choose classical smoothness conditions, but rather formulate the regularity in terms of the total variation norm of the mapping $t \mapsto G_t$, measured in $L_2(P)$. In particular, we suppose that for some $\tilde{\Gamma} > 0$,

$$\sum_{t=2}^n (\mathbb{E}|G_{t,l}(\epsilon_0) - G_{t-1,l}(\epsilon_0)|^q)^{\frac{1}{q}} \leq \tilde{\Gamma} \cdot \tilde{\Theta}, \quad l = 1, \dots, d. \tag{P.2}$$

A few results require to assume

$$X_t \text{ has finite and stationary eighth moment and } \beta > 2. \tag{P.3}$$

We highlight that (P.1) and (P.2) only impose conditions on the individual components $l = 1, \dots, d$, and are thus rather easy to verify even for high-dimensional models. To emphasize, we impose no assumption on the dependency among the d components.

Now consider the projection of the d -variate time series X_t to a m -dimensional time series Z_t , with $m \ll d$. That is, for a matrix $V \in \mathbb{R}^{m \times d}$, we consider the

time series $Z_t = VX_t$. We measure the size of the matrix V by the operator norm, $\|V\|_\infty$, with respect to the maximum vector norm, i.e.

$$\|V\|_\infty = \sup_{w \in \mathbb{R}^d} \frac{\|Vw\|_\infty}{\|w\|_\infty} = \max_{l=1, \dots, m} \|V_{l, \cdot}\|_1,$$

where $V_{l, \cdot}$ denotes the l -th row of the matrix V . This choice of norm uniformly controls the ℓ_1 -norms of the projection vectors and is motivated by high-dimensional statistical problems. Indeed, as well known, if d is large compared to n , classical statistical methods usually fail, because even strong signals are overlaid by too many random noise sources (noise accumulation) and spurious correlations occur. To overcome the curse of high dimensionality, sparse methods typically use $\|\cdot\|_r$ -norms with $1 \leq r < 2$, often leading to s -sparse solutions, i.e. with only s active (non-zero) coordinates. For example, sparse PCA constructs s -sparse directions u on which the data vectors are projected by calculating $u^T X_t$, and the variance of such a coefficients $u^T X_t$, given by the quadratic form $u^T \Sigma_n u$ and coinciding with the associated eigenvalue when u is an eigenvector, provides information about the importance of that direction and is thus of interest. Similarly, when predicting a future response \tilde{Y}_t by a lasso regression on X_t , one calculates $\hat{\beta}_n^T \tilde{X}_t$ for future regressors \tilde{X}_t , where the lasso estimate $\hat{\beta}_n$ is s -sparse and thus selects $s \ll d$ variables, so that it behaves like a low-dimensional statistic circumventing the issues arising for large dimension. The variance $\hat{\beta}_n^T \Sigma_n \hat{\beta}_n^T$ of the prediction $\hat{\beta}_n^T \tilde{X}_t$ is a natural measure of the prediction accuracy. When considering slightly more general vectors v with bounded $\|v\|_1$ -norm (uniformly in d and n), then we have the bound $|v^T X_t| \leq \|v\|_1 \max_{1 \leq j \leq d} |X_{tj}|$ at our disposal. Indeed, such projections have bounded moments under weak assumptions guaranteed by (P.1) as summarized in the following lemma, which no longer holds if $\|v\|_2$ is uniformly bounded but $\|v\|_1 \rightarrow \infty$.

Lemma 2.1. *Let $q \geq 1$ and suppose that $\max_{1 \leq j \leq d} \mathbb{E}|X_{tj}|^q \leq \tilde{\Theta}$ for some constant C and all d, n , cf. (P.1). Then for any $v \in \mathbb{R}^s$ with norm $\|v\|_1$ uniformly bounded by some M*

$$(\mathbb{E}|v^T X_t|^q)^{\frac{1}{q}} \leq \|v\|_1 \max_{1 \leq j \leq d} (\mathbb{E}|X_{tj}|^q)^{\frac{1}{q}} \leq \tilde{\Theta} M$$

The formulation of our first main result requires the two rates

$$\chi(q, \beta) = \begin{cases} \frac{q-2}{6q-4}, & \beta \geq \frac{3}{2}, \\ \frac{(\beta-1)(q-2)}{q(4\beta-3)-2}, & \beta \in (1, \frac{3}{2}). \end{cases}, \quad \xi(q, \beta) = \begin{cases} \frac{q-2}{6q-4}, & \beta \geq 3, \\ \frac{(\beta-2)(q-2)}{(4\beta-6)q-4}, & \frac{3+\frac{2}{q}}{1+\frac{2}{q}} < \beta < 3, \\ \frac{1}{2} - \frac{1}{\beta}, & 2 < \beta \leq \frac{3+\frac{2}{q}}{1+\frac{2}{q}}. \end{cases}$$

Theorem 2.2. *Let $X_t = G_t(\epsilon_t)$ with $\mathbb{E}(X_t) = 0$ be such that (P.1) holds for some $q > 2$, $\beta > 1$, and let $Z_t = VX_t$ for some $V \in \mathbb{R}^{m \times d}$ such that $m \leq cn$ for some $c > 0$. Then, on a potentially different probability space, there*

exist random vectors $(X'_t)_{t=1}^n \stackrel{d}{=} (X_t)_{t=1}^n$ and independent, mean zero, Gaussian random vectors Y'_t such that

$$\left(\mathbb{E} \max_{k \leq n} \left\| \frac{1}{\sqrt{n}} \sum_{t=1}^k (V X'_t - Y'_t) \right\|_2^2 \right)^{\frac{1}{2}} \leq C \tilde{\Theta} \|V\|_\infty \sqrt{m \log(n)} \left(\frac{m}{n}\right)^{\chi(q,\beta)}. \quad (2)$$

for some universal constant $C = C(q, \beta, c)$.

If $\beta > 2$, the local long-run variance $\Xi_t(v) = \sum_{h=-\infty}^\infty V \text{Cov}(G_t(\epsilon_0), G_t(\epsilon_h)) V^T$ is well defined. If (P.2) is satisfied as well, then there exist random vectors $(X'_t)_{t=1}^n \stackrel{d}{=} (X_t)_{t=1}^n$ and independent, mean zero, Gaussian random variables $Y'_t \sim \mathcal{N}(0, \Xi_t(v))$ such that

$$\left(\mathbb{E} \max_{k \leq n} \left\| \frac{1}{\sqrt{n}} \sum_{t=1}^k (V X'_t - Y'_t) \right\|_2^2 \right)^{\frac{1}{2}} \leq C \tilde{\Theta} \tilde{\Gamma}^{\frac{1}{2} \frac{\beta-2}{\beta-1}} \|V\|_\infty \sqrt{m \log(n)} \left(\frac{m}{n}\right)^{\xi(q,\beta)}. \quad (3)$$

The rate of the approximation (2) is faster than (3). The difference is that in the first case, the covariance structure of the approximating Gaussian random vectors Y'_t is not explicit. In the second case, the distribution of the approximating random vectors is explicit, at the price of assuming some temporal regularity of the stochastic process X_t . In Section 5, we will describe a bootstrap scheme to perform inference based on Theorem 2.2.

Neglecting the high-dimensional context for a moment, the projection $Z_t = V X_t$ may be analyzed as a multivariate, nonstationary time series of fixed dimension m . In this setting, Wu and Zhou (2011) and Karmakar and Wu (2020) study sequential Gaussian approximations, the latter obtaining optimal rates in n which are faster than the rates of Theorem 2.2. However, they require a lower bound on the covariance matrices of Z_t , while our result does not need this assumption. This is particularly relevant because we do not impose any conditions on the dependency among the components of X_t . As an extreme example, if all d components are independent, then a l_1 -bounded projection will lead to some concentration due to averaging, such that a lower bound on the variance of the projection can in general not be guaranteed. This prevents an application of the result of Karmakar and Wu (2020) in the present situation. Moreover, the covariance of the approximating random vectors in Karmakar and Wu (2020) is not explicit. Instead, we employ a result on Gaussian couplings presented in Mies and Steland (2023), which is in turn enabled by a recent result of Eldan, Mikulincer and Zhai (2020).

3. Asymptotics of high-dimensional shrinkage covariance matrix estimators and their projections

For a d -variate centered time series X_t as above, denote the sample covariance matrix by

$$\hat{\Sigma}_{n,n} = \frac{1}{n} \sum_{t=1}^n X_t X_t^T,$$

i.e. the empirical lag-0 cross-covariance matrix of the time series X_t . The respective partial sums are denoted by $\hat{\Sigma}_{k,n} = \frac{1}{n} \sum_{t=1}^k X_t X_t^T$, $k = 1, \dots, n$, and its mean by

$$\Sigma_n = \mathbb{E} \hat{\Sigma}_{n,n} = \frac{1}{n} \sum_{t=1}^n \text{Cov}(X_t).$$

Hence, we account for nonstationarity of the time series X_t by averaging the covariance matrices at all times $t = 1, \dots, n$.

The central observation enabling our subsequent analysis is that the matrix-valued time series $X_t X_t^T$ fits within the framework of Section 2.

Proposition 3.1. *If the time series $X_t = G_t(\epsilon_t)$ satisfies (P.1) resp. (P.2) with power $2q$ and factor Θ , then the d^2 -variate time series $\tilde{X}_t = X_t X_t^T = \tilde{G}_t(\epsilon_t)$ satisfies (P.1) resp. (P.2) with power q and factor Θ^2 .*

It has been suggested by [Steland and von Sachs \(2017\)](#) to perform inference on the covariance matrix via the quadratic form $v^T \hat{\Sigma}_{n,n} v$ for some vector $v \in \mathbb{R}^d$, i.e., the variance of the projection $v^T X_t$. As already discussed above, such projections are ubiquitous in statistical problems, for example arising in optimal portfolio selection, as linear predictors in regression models, or when reducing dimensionality by principal component analysis. The associated empirical version, $v^T \hat{\Sigma}_{n,n} v = \frac{1}{n} \sum_{t=1}^n |v^T X_t|^2$, may be interpreted as the estimated variance of the univariate projection $\tilde{X}_t = v^T X_t$. The results of [Steland and von Sachs \(2017\)](#) and [Steland and von Sachs \(2018\)](#) on Gaussian approximations of $v^T \hat{\Sigma}_n v$, related change-point procedures, and Shrinkage estimators, impose the assumption that the process X_t is a linear time series with a single innovation process, which may be interpreted as a special type of single-factor model, although approximate vector autoregressions and spiked covariance models are included under regularity conditions, see [Steland \(2020\)](#). Extensions to multi-factor models with infinitely many factors and general multivariate linear processes with finite-dimensional innovations are studied in [Bours and Steland \(2021\)](#). Although such processes, especially factor models, provide good description for many application scenarios and are widely used, they nevertheless restrict the dependence structure of the d component processes and rule out phenomena such as conditional heteroscedasticity and nonlinearity. Here, by relying on the results of the previous section, we allow for nonstationary nonlinear time series and impose much weaker assumptions on the dependence of the components.

To estimate Σ_n and related quadratic forms $v^T \Sigma_n v$, we shall shrink the nonparametric estimator $\hat{\Sigma}_{n,n}$ towards structured targets, in order to overcome

various issues of the sample covariance matrix arising if $d \gg n$. A covariance matrix Σ is called bandable if $\Sigma_{l,l'} \leq g(|l - l'|)$ for some function $g(x) \rightarrow 0$ as $x \rightarrow \infty$. In this situation, optimal rates of convergence can be achieved via thresholding, see (Cai, Ren and Zhou, 2016, Thm. 7) and references therein. In particular, if $g(x) \propto |x|^{-\alpha-1}$ for some $\alpha > 0$, then a rate-optimal estimator is given by the tapering estimator

$$\hat{\Sigma}_{k,n}^\dagger = ((\hat{\Sigma}_{k,n})_{l,l'} \omega(|l - l'|))_{l,l'=1}^d,$$

$$\omega(x) = \begin{cases} 1, & x \leq \frac{\tau}{2}, \\ 2 - \frac{2x}{\tau}, & \frac{\tau}{2} < x \leq \tau, \\ 0, & x > \tau, \end{cases}$$

The optimal rate is achieved by the threshold $\tau = \min(n^{\frac{1}{2\alpha+c}}, d)$ with $c = 2$ under the Frobenius norm $\|\cdot\|_F$ and with $c = 1$ under the spectral norm $\|\cdot\|$, see Cai, Zhang and Zhou (2010). In particular, the optimal rates are

$$d^{-1} \|\hat{\Sigma}_{k,n}^\dagger - \Sigma\|_F^2 = \mathcal{O}\left(n^{-\frac{2\alpha+1}{2\alpha+2}} + \frac{d}{n}\right), \quad \|\hat{\Sigma}_{k,n}^\dagger - \Sigma\|^2 = \mathcal{O}\left(n^{-\frac{2\alpha}{2\alpha+1}} + \frac{\log d}{n}\right).$$

Another approach is to impose a Toeplitz structure, i.e. $\Sigma_{l,l'} = \sigma_{|l-l'|}$, for some sequence $\sigma_0, \sigma_1, \dots$. The corresponding tapered Toeplitz estimator of Σ is

$$\hat{\Sigma}_{k,n}^\circ = (\hat{\sigma}_{|l-l'|} \omega(|l - l'|))_{l,l'=1}^d,$$

$$\hat{\sigma}_m = \frac{1}{d-m} \sum_{l-l'=m} (\hat{\Sigma}_{k,n})_{l,l'}, \quad 0 \leq m < d,$$

with tapering weights $\omega(x)$ as above. If $|\sigma_m| \leq Cm^{-\alpha-1}$, then the optimal rate under the spectral norm is achieved by the threshold choice $\tau = (nd/\log(nd))^{\frac{1}{2\alpha+1}}$, see (Cai, Ren and Zhou, 2016, Thm. 9) and Cai, Ren and Zhou (2013).

The following theorem shows that the full-sample estimators $\hat{\Sigma}_n^\dagger = \hat{\Sigma}_{n,n}^\dagger$ and $\hat{\Sigma}_n^\circ = \hat{\Sigma}_{n,n}^\circ$ are consistent in the (rescaled) Frobenius norm, under mild regularity conditions. Moreover, $\hat{\Sigma}_n^\dagger$ is rate-optimal for bandable covariance matrices. To the best of our knowledge, these estimators have not yet been studied under such a general nonstationary nonlinear time series model as given by (P.1) and (P.2). Indeed, the optimal rates and thresholds have been obtained under the assumption that X_t are iid random vectors with mean zero and a bandable or Toeplitz covariance matrix Σ satisfying further regularity conditions.

Denote by Σ_n^\dagger and Σ_n° the theoretical oracle estimators associated to the estimators $\hat{\Sigma}_{n,n}^\dagger$ and $\hat{\Sigma}_{n,n}^\circ$, respectively, obtained by replacing in their definitions the estimates $(\hat{\Sigma}_{n,n})_{i,j}$ by the true covariances $(\Sigma_{n,n})_{i,j}$. That is, $\Sigma_n^\circ = \mathbb{E}(\hat{\Sigma}_n^\circ)$ and $\Sigma_n^\dagger = \mathbb{E}(\hat{\Sigma}_n^\dagger)$. Recall also that $\Sigma_n = \mathbb{E}(\hat{\Sigma}_{n,n})$. Denote for $A, B \geq 0$ the scaled inner product by $\langle A, B \rangle_* = A \cdot B/d$ and let $\|A\|_{F_*} = \sqrt{\langle A, A \rangle_*}$ be the associated scaled Frobenius matrix norm. This scaling ensures that the unit matrix has norm 1 whatever the dimension d . The following theorem studies

the above estimators under the general time series model given by (P.1) and (P.2) and provides sufficient conditions on the growth of dimension d relative to the sample size when measuring the estimation error in terms of the scaled Frobenius risk.

Theorem 3.2. *Suppose that $\{X_t\}$ satisfies (P.1) with $q \geq 4$ and $\beta > 1$, and that $\text{Cov}(X_t)$ is bandable with exponent $\alpha > 0$ uniformly for all t , i.e.*

$$\text{Cov}(X_t)_{i,j} \leq C|i-j|^{-\alpha-1}.$$

(i) *For any threshold τ , it holds*

$$\mathbb{E}\|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2 \leq C \left\{ \tau^{-2\alpha-1} + C \frac{\tilde{\Theta}^4}{(\beta-1)^2} \frac{\min(\tau, d)}{n} \right\}.$$

Choosing $\tau = \min(n^{\frac{1}{2\alpha+2}}, d)$, it holds

$$\mathbb{E}\|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2 \leq C(\Theta, \beta) \min(n^{-\frac{2\alpha+1}{2\alpha+2}}, \frac{d}{n}).$$

(ii) *If Σ_n is a Toeplitz matrix, then*

$$\mathbb{E}\|\hat{\Sigma}_n^\diamond - \Sigma_n\|_{F^*}^2 \leq C \left\{ \tau^{-2\alpha-1} + \frac{\tilde{\Theta}^4}{(\beta-1)^2} \frac{\min(\tau, d)}{n} \right\}.$$

Choosing $\tau = \min(n^{\frac{1}{2\alpha+2}}, d)$, it holds

$$\mathbb{E}\|\hat{\Sigma}_n^\diamond - \Sigma_n\|_{F^*}^2 \leq C(\Theta, \beta) \min(n^{-\frac{2\alpha+1}{2\alpha+2}}, \frac{d}{n}).$$

In particular, the optimal rate of convergence of the tapering estimator $\hat{\Sigma}_n^\dagger$ carries over to the dependent, nonstationary case, with identical threshold values. This means, combining Theorem 3.2 and Theorem 4 of Cai, Zhang and Zhou (2010) yields the minimax result

$$\inf_{\hat{\Sigma}_n} \sup_{\mathcal{G}_\alpha} \mathbb{E}\|\hat{\Sigma}_n - \Sigma_n\|_{F^*} \asymp \min(n^{-\frac{2\alpha+1}{2\alpha+2}}, \frac{d}{n}),$$

where the optimal rate is attained by the tapering estimator $\hat{\Sigma}_n^\dagger$ when using the threshold $\tau = \min(n^{\frac{1}{2\alpha+2}}, d)$. Here, \mathcal{G}_α denotes the class of distributions of $\mathbf{X}_n = (X_1, \dots, X_n)^\top$ such that model (1) holds under assumptions (P.1)-(P.3), Σ_n is bandable with decay parameter α , and the largest eigenvalue of Σ_n is bounded by some given constant.

Note that the derived error bounds for the tapering estimator and the Toeplitz estimator are identical; in particular, the pooling of covariances via $\hat{\sigma}_m$ does not give an improved error bound, which is a consequence of our weak assumptions. As we do not impose any assumptions on the dependency between the components, the averaging occurring in the definition of $\hat{\sigma}_m$ can not be leveraged for tighter controls on the estimation error. Moreover, we note that the rate for $\hat{\Sigma}_n^\diamond$ is worse compared to the results of Cai, Ren and Zhou (2013). However, our bound

is in terms of the Frobenius-star norm instead of the operator norm. Furthermore, the proof of Cai, Ren and Zhou (2013) crucially relies on an exponential inequality for spectral density estimators, which is available for iid samples but not under our general nonstationary time series model. Theorem 3.2 ensures that under our assumptions the estimators used to construct the target are consistent at a good convergence rate. This suggests to put $\tau = \min(n^{\frac{1}{2\alpha+2}}, d)$.

A shrinkage estimator may be defined for any weight $w \in \Delta = \{(w_1, w_2, w_3) \in [0, 1]^3 : w_1 + w_2 + w_3 = 1\}$, as

$$\hat{\Sigma}_{k,n}^w = w_1 \hat{\Sigma}_{k,n} + w_2 \hat{\Sigma}_{k,n}^\dagger + w_3 \hat{\Sigma}_{k,n}^\diamond.$$

It turns out that the theory presented in Section 2 may be applied to the quadratic form $v^T \hat{\Sigma}_n^w v$ for any shrinkage weight w , and hence also for the individual statistics $v^T \hat{\Sigma}_n v$, $v^T \hat{\Sigma}_n^\dagger v$, and $v^T \hat{\Sigma}_n^\diamond v$. The crucial observation is that these quadratic forms can be regarded as linear projections of the $(d \times d)$ -variate time series $X_t X_t^T$ with respect to the Frobenius inner product $A \cdot B = \sum_{i,j} a_{ij} b_{ij}$ defined for $d \times d$ matrices $A = (a_{ij})_{i,j}$ and $B = (b_{ij})_{i,j}$, which coincides with the usual inner product on \mathbb{R}^{d^2} of the vectorized matrices. Indeed, we have the representations

$$v^T \hat{\Sigma}_{k,n}^\dagger v = \nu^2 \cdot \hat{\Sigma}_{k,n}, \quad v^T \hat{\Sigma}_{k,n}^\diamond v = \nu^3 \cdot \hat{\Sigma}_{k,n},$$

for the projection weighting matrices $\nu^1, \nu^2, \nu^3 \in \mathbb{R}^{d \times d} \equiv \mathbb{R}^{d^2}$ given by

$$\begin{aligned} \nu^1 &= (v_l v_{l'})_{l,l'=1}^d, & \|\nu^1\|_1 &= \sum_{l,l'=1}^d |v_l v_{l'}| = \|v\|_1^2, \\ \nu^2 &= (v_l v_{l'} \omega(|l-l'|))_{l,l'=1}^d, & \|\nu^2\|_1 &\leq \|\nu^1\|_1 = \|v\|_1^2, \\ \nu^3 &= \left(\frac{\omega(|l-l'|)}{d-|l-l'|} \sum_{i=j=l-l'}^d v_i v_j \right)_{l,l'=1}^d, & \|\nu^3\|_1 &\leq \|\nu^1\|_1 = \|v\|_1^2 \end{aligned}$$

In particular, we may write $v^T X_t X_t^T v = \nu^1 \cdot (X_t X_t^T)$ as the Frobenius inner product.

With Theorem 2.2, we obtain the following result.

Theorem 3.3. *Suppose that the time series $X_t = G_t(\epsilon_t)$ satisfies (P.1) and (P.2) with power $q > 4$ and exponent $\beta > 2$. Then, the time series X_t may be defined on a different probability space, such that there exist independent, mean zero, Gaussian random vectors $\eta_t \in \mathbb{R}^3$, such that*

$$\begin{aligned} & \left(\mathbb{E} \max_{k=1, \dots, n} \max_{w \in \Delta} \left| \sqrt{n} v^T \left[\hat{\Sigma}_{k,n}^w - \mathbb{E} \hat{\Sigma}_{k,n}^w \right] v - w^T \frac{1}{\sqrt{n}} \sum_{t=1}^k \eta_t \right|^2 \right)^{\frac{1}{2}} \\ & \leq C \tilde{\Theta}^2 \tilde{\Gamma}^{\frac{1}{2} \frac{\beta-1}{\beta-2}} \|v\|_1^2 \sqrt{\log(n)} n^{-\xi(q,\beta)}, \end{aligned}$$

for some universal constant $C = C(q, \beta)$. The covariance of the Gaussian vectors $\eta_t \sim \mathcal{N}(0, \Xi_t)$ is given by

$$(\Xi_t)_{i,j} = \sum_{h=-\infty}^{\infty} \text{Cov}(\nu^i \cdot G_t(\epsilon_0)G_t(\epsilon_0)^T, \nu^j \cdot G_t(\epsilon_h)G_t(\epsilon_h)^T), \quad i, j \in \{1, 2, 3\}.$$

It is worth mentioning that the bound in Theorem 3.3 allows that $\|v\|_1$ increases as the dimension, d , or the sample size, n , grow, as long as $\|v\|_1 = o((\log n)^{1/4}n^{-\xi(q,\beta)/2})$.

For simplicity of presentation, Theorem 3.3 is formulated for a single projection vector v . However, we might as well use multiple vectors $v^{(1)}, \dots, v^{(m)} \in \mathbb{R}^d$ for the bilinear forms, and obtain a multivariate Gaussian approximation. Denote by $\nu_{(j)}^1, \nu_{(j)}^2, \nu_{(j)}^3$ the projection weighting matrices based on the vector $v^{(j)}$, $j = 1, \dots, m$.

Theorem 3.4. *Suppose that the time series $X_t = G_t(\epsilon_t)$ satisfies (P.1) resp. (P.2) with power $q > 4$ and exponent $\beta > 2$, and suppose that $m \leq cn$ for some constant c . Then, the time series X_t may be defined on a different probability space, such that there exist independent, mean zero, Gaussian random vectors $\eta_t^{(s)} \in \mathbb{R}^3$, for $s = 1, \dots, m$ and $t = 1, \dots, n$, independent for different t , such that*

$$\begin{aligned} & \left(\mathbb{E} \max_{k=1, \dots, n} \max_{w \in \Delta} \sum_{s=1}^m \left| \sqrt{n}(v^{(s)})^T \left[\hat{\Sigma}_{k,n}^w - \mathbb{E} \hat{\Sigma}_{k,n}^w \right] v^{(s)} - w^T \frac{1}{\sqrt{n}} \sum_{t=1}^k \eta_t^{(s)} \right|^2 \right)^{\frac{1}{2}} \\ & \leq C \tilde{\Theta}^2 \tilde{\Gamma}^{\frac{1}{2}} \frac{\beta-1}{\beta-2} \left(\max_{s=1, \dots, m} \|v^{(s)}\|_1^2 \right) \sqrt{m \log(n)} \left(\frac{m}{n}\right)^{\xi(q,\beta)}, \end{aligned}$$

for some universal constant $C = C(q, \beta, c)$. The covariance of the Gaussian vectors is given by

$$\begin{aligned} & \text{Cov}((\eta_t^{(s)})_i, (\eta_t^{(s')})_i) = (\Xi_t^{(s,s')})_{i,j} \\ & = \sum_{h=-\infty}^{\infty} \text{Cov}(\nu_{(s)}^i \cdot G_t(\epsilon_0)G_t(\epsilon_0)^T, \nu_{(s')}^j \cdot G_t(\epsilon_h)G_t(\epsilon_h)^T), \quad i, j \in \{1, 2, 3\}. \end{aligned}$$

Remark 1. It is also possible to choose the vector v randomly. In this situation, Theorems 3.3 and 3.4 are still applicable conditionally on v , provided that the choice of v is independent from the data. Moreover, v may also be chosen in a data-driven way. This is, for instance, relevant for financial risk where $v = \hat{v}$ represents the optimal portfolio weights according to some criterion as discussed in the introduction. For our results to be applicable, it is important that the weights \hat{v} are based on historical data, say from time $t = -B$ and older. For sufficiently large B , assumption P.1 ensures that \hat{v} and $\hat{\Sigma}_{k,n}$ are essentially independent. The same applies to the choice of the shrinkage weights w , which may be determined deterministically, randomly, or based on historical data.

4. Optimal shrinkage weights

The question arises, how the shrinking weights relate to the theoretical (oracle) performance of the resulting shrinkage estimator. To pursue such a study, we first consider the problem to combine the nonparametric estimator with oracles of the two targets within a framework proposed by [Ledoit and Wolf \(2004\)](#) and aim at minimizing the squared Frobenius risk. This leads to a convex but box-constrained optimization problem allowing for an explicit interior solution (if it exists), which we briefly discuss. We propose consistent estimators of the unknown quantities determining the optimization problem and its solution(s) leading to data-adaptive bona fide estimates. A similar approach has been pursued by [Bodnar, Gupta and Parolya \(2014\)](#), who derive optimal shrinkage weights for a general shrinkage target and construct bona fide estimators. In contrast, we consider not one but two shrinkage targets, which are furthermore data-dependent, and our bona-fide shrinkage weights need to account for temporal dependence.

For simplicity of presentation, we elaborate this for the full sample of size n ; the corresponding sequential estimates and oracles based on the first k observations can then be derived easily. Write $\hat{\Sigma}_n = \hat{\Sigma}_{n,n}$ and recall that Σ_n^\dagger and Σ_n° are the theoretical oracle estimators associated to the estimators $\hat{\Sigma}_{n,n}^\dagger$ and $\hat{\Sigma}_{n,n}^\circ$, respectively, obtained by replacing in their definitions the estimates $(\hat{\Sigma}_{n,n})_{i,j}$ by the true covariances $(\Sigma_{n,n})_{i,j}$. Equivalently, $\Sigma_n^\circ = \mathbb{E}\hat{\Sigma}_n^\circ$ and $\Sigma_n^\dagger = \mathbb{E}\hat{\Sigma}_n^\dagger$.

We wish to determine the optimal ideal shrinking weights such that the resulting (scaled) Frobenius risk is minimized when combining the nonparametric estimator $\hat{\Sigma}_n$ with the oracle estimators Σ_n^\dagger and Σ_n° . The scaled mean squared error of the sample covariance matrix is given by $\text{MSE}(\hat{\Sigma}_n) = \text{MSE}(\hat{\Sigma}_n; \Sigma_n) = E\|\hat{\Sigma}_n - \Sigma_n\|_{F^*}^2$, and we consider the risk minimization problem

$$\begin{aligned} \min_w \quad & \mathbb{E}\|\Sigma_n^w - \Sigma_n\|_{F^*}^2 \\ \text{where} \quad & \Sigma_n^w = w_1\hat{\Sigma}_n + w_2\Sigma_n^\dagger + w_3\Sigma_n^\circ, \quad w_1 + w_2 + w_3 = 1. \end{aligned} \tag{4}$$

Further, let

$$E_n^\dagger = \|\Sigma_n^\dagger - \Sigma_n\|_{F^*}^2, \quad E_n^\circ = \|\Sigma_n^\circ - \Sigma_n\|_{F^*}^2$$

be the squared approximation errors corresponding to the oracles, and

$$D_n = \langle \Sigma_n^\dagger - \Sigma_n, \Sigma_n^\circ - \Sigma_n \rangle_*$$

Substituting $w_1 = 1 - w_2 - w_3$ and observing that $\mathbb{E}\langle \hat{\Sigma}_n - \Sigma_n, A \rangle = 0$, for $A \in \{\Sigma_n^\dagger, \Sigma_n^\circ\}$, it suffices to determine minimizers of

$$f(w_2, w_3) = (1 - w_2 - w_3)^2 \mathbb{E}\|\hat{\Sigma}_n - \Sigma_n\|_{F^*}^2 + w_2^2 E_n^\dagger + w_3^2 E_n^\circ + 2w_2 w_3 D_n. \tag{5}$$

Adding the constraints $0 \leq w_1, w_2 \leq 1$ leads to a quadratic programming problem under box constraints, which generally requires numerical algorithms for its solution. We discuss the special case of an interior solution, where explicit formulas can be derived, at the end of this section. The optimal weights depend

on the unknown values $MSE(\hat{\Sigma}_n) = \mathbb{E}\|\hat{\Sigma}_n - \Sigma_n\|_{F^*}^2$, E_n^\dagger , E_n° , and D_n . To obtain a data-driven solution for the optimal weights, we need consistent estimators of these quantities, which are also interesting in their own right.

Observe that

$$MSE(\hat{\Sigma}_n) = \frac{1}{nd} \sum_{i,j=1}^d \text{Var}(\sqrt{n}(\hat{\Sigma}_n)_{i,j}).$$

Under the assumptions of Proposition 3.1 $\sqrt{n}(\hat{\Sigma}_n)_{i,j} = \frac{1}{\sqrt{n}} \sum_{t=1}^n (X_t)_i (X_t)_j$ can be approximated in L_2 by $\frac{1}{\sqrt{n}} \sum_{t=1}^n Y_{t,i,j}$ for independent Gaussian random variables

$$Y_{t,i,j} \sim N(0, \xi_{t,i,j}^2) \quad \text{with} \quad \xi_{t,i,j}^2 = \sum_{h=-\infty}^{\infty} \text{Cov}(G_{t,i}(\epsilon_0)G_{t,j}(\epsilon_0), G_{t,i}(\epsilon_h)G_{t,j}(\epsilon_h)).$$

This yields

$$MSE(\hat{\Sigma}_n) = \frac{1}{nd} \sum_{i,j=1}^d \frac{1}{n} \sum_{t=1}^n \xi_{t,i,j}^2 + o(1).$$

Sancetta (2008) proposed an estimator of $MSE(\hat{\Sigma}_n)$ by estimating the long-run-variance parameters $\text{Var}(\sqrt{n}(\hat{\Sigma}_n)_{i,j})$ by the long-run-variance estimators

$$\widehat{\text{Var}}(\sqrt{n}(\hat{\Sigma}_n)_{i,j}) = \hat{\Gamma}_{n,i,j}(0) + 2 \sum_{s=1}^{n-1} K(s/b) \hat{\Gamma}_{n,i,j}(s),$$

where K is a positive weighting function, $b > 0$ is a bandwidth, and

$$\hat{\Gamma}_{n,i,j}(s) = \frac{1}{n} \sum_{t=1}^{n-s} [(X_t)_i (X_t)_j - (\hat{\Sigma}_n)_{i,j}] \cdot [(X_{t+s})_i (X_{t+s})_j - (\hat{\Sigma}_n)_{i,j}].$$

The mean square error may then be estimated as

$$\widehat{MSE}(\hat{\Sigma}_n) = \frac{1}{nd} \sum_{i,j=1}^d \widehat{\text{Var}}(\sqrt{n}(\hat{\Sigma}_n)_{i,j}).$$

The result is as follows, see (Sancetta, 2008, Lemma 1).

Lemma 4.1. *If (P.1)-(P.3) hold, $K : \mathbb{R} \rightarrow \mathbb{R}$ is a decreasing positive càdlàg-function with $\lim_{x \rightarrow 0^+} K(x) = 1$ and $\int_0^\infty K^2(x)dx < \infty$, and $b \rightarrow \infty$ with $b = o(\sqrt{n})$, then*

$$\left| \widehat{MSE}(\hat{\Sigma}_n) - MSE(\hat{\Sigma}_n) \right| = o\left(\frac{d}{n}\right).$$

In practice, the bandwidth b needs to be chosen in a data-driven way, e.g. by the method described in Andrews (1991). As we essentially estimate the long run variance of many scalar entries, the latter method is expected to perform well even in the present high-dimensional context.

Let us now derive consistent estimators for the approximation errors E_n^\dagger and E_n° and the inner product D_n . The quantities $\Delta_n^\circ = \Sigma_n - \hat{\Sigma}_n^\circ$ and $E_n^\circ = \|\Sigma_n - \hat{\Sigma}_n^\circ\|_{F^*}^2$ can be estimated by plugging in the estimators $\hat{\Sigma}_n^\circ$ and $\hat{\Sigma}_n^\dagger$. Hence, define

$$\hat{\Delta}_n^\circ = \hat{\Sigma}_n^\dagger - \hat{\Sigma}_n^\circ, \quad \hat{E}_n^\circ = \|\hat{\Sigma}_n^\dagger - \hat{\Sigma}_n^\circ\|_{F^*}^2.$$

To estimate E_n^\dagger observe that $\Delta_n^\dagger = (\Delta_{n,ij}^\dagger)_{i,j} = \Sigma_n - \Sigma_n^\dagger$ is given by

$$\Delta_n^\dagger = \begin{cases} 0, & |i-j| \leq \tau/2, \\ (1 - \omega(|i-j|))\Sigma_{n,ij}, & \tau/2 < |i-j| \leq \tau, \\ \Sigma_{n,ij}, & |i-j| > \tau. \end{cases}$$

Estimation of Δ_n^\dagger is delicate, since this involves estimating $\sum_{m>\tau} \sum_{|i-j|=m} 1 = (d-\tau)(d-\tau+1)/2 = O(d^2)$ off-diagonal terms which would lead to the unsatisfactory condition $d^2 = o(n)$. For the class of bandable covariance matrices under consideration, we can proceed by introducing a second threshold and estimate only the covariances on the diagonals $\tau/2 + 1, \dots, \sigma$. This means, we estimate Δ_n^\dagger by $\hat{\Delta}_n^\dagger = (\hat{\Delta}_{n,ij}^\dagger)_{i,j}$ with

$$\hat{\Delta}_{n,ij}^\dagger = \begin{cases} 0, & |i-j| \leq \tau/2, \\ (1 - \omega(|i-j|))\hat{\Sigma}_{n,ij}, & \tau/2 < |i-j| \leq \tau, \\ \hat{\Sigma}_{n,ij}, & \tau < |i-j| \leq \sigma, \\ 0, & |i-j| > \sigma, \end{cases}$$

for some threshold σ with $\tau \ll \sigma \ll d$, and put

$$\hat{E}_n^\dagger = \|\hat{\Delta}_n^\dagger\|_{F^*}^2.$$

Lastly, the estimation of $D_n = \langle \Sigma_n - \Sigma_n^\dagger, \Sigma_n - \Sigma_n^\circ \rangle_*$ can be based on the above estimators of $\Delta_n^\dagger = \Sigma_n - \Sigma_n^\dagger$ and $\Delta_n^\circ = \Sigma_n - \Sigma_n^\circ$,

$$\hat{D}_n = \langle \hat{\Delta}_n^\dagger, \hat{\Delta}_n^\circ \rangle_*.$$

The following results, which provide L_1 - and L_2 -bounds, respectively, ensuring consistency, treat the highdimensional case $d \geq n^{\frac{1}{2\alpha+2}}$.

Theorem 4.2. *Suppose that $\{X_t\}$ satisfies (P.1) and (P.2) and $\text{Var}(X_t)$ is bandable for all t . Further, assume that $d \geq n^{\frac{1}{2\alpha+2}}$, and the estimators $\hat{\Sigma}_n^\dagger$, $\hat{\Delta}_n^\dagger$ and $\hat{\Sigma}_n^\circ$ are calculated with thresholds $\tau_n^\dagger = n^{\frac{1}{2\alpha+2}}$, $\tau_n^\circ = (nd/\log(nd))^{\frac{1}{2\alpha+1}}$, and $\sigma_n^\dagger = n^{\frac{1+s}{2\alpha+2}}$ for some $0 < s < 2\alpha$.*

(i) *It holds*

$$\begin{aligned} \mathbb{E}\|\hat{\Delta}_n^\circ - \Delta_n^\circ\|_{F^*} &= \mathcal{O}\left(n^{-\frac{\alpha+1/2}{2\alpha+2}}\right), \\ \mathbb{E}\|\hat{\Delta}_n^\circ - \Delta_n^\circ\|_{F^*}^2 &= \mathcal{O}\left(n^{-\frac{2\alpha+1}{2\alpha+2}}\right). \end{aligned}$$

and

$$\mathbb{E}|\hat{E}_n^\circ - E_n^\circ| = \mathbb{E}\left|\|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger\|_{F^*}^2 - \|\Sigma_n^\circ - \Sigma_n\|_{F^*}^2\right| = \mathcal{O}\left(n^{-\frac{\alpha+1/2}{2\alpha+2}}\right).$$

(ii) We have

$$\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2 = \mathcal{O}\left(n^{-\frac{2\alpha+1-s}{2\alpha+2}}\right)$$

and

$$\mathbb{E}|\hat{E}_n^\dagger - E_n^\dagger| = \mathbb{E}\left|\|\hat{\Delta}_n^\dagger\|_{F^*} - \|\Delta_n^\dagger\|_{F^*}\right| = \mathcal{O}\left(n^{-\frac{2\alpha+1-s}{2(2\alpha+2)}}\right).$$

(iii) It holds

$$\mathbb{E}|\hat{D}_n - D_n| = \mathcal{O}\left(\sqrt{\left(\frac{d}{\log(nd)}\right)^{\frac{1}{2\alpha+1}} n^{-\frac{2\alpha-1-s}{2\alpha+1}}}\right),$$

and if $d = \mathcal{O}(n^\kappa)$ for some $\kappa > 0$, then a sufficient condition for L_1 -consistency of \hat{D}_n is $\kappa \leq 2\alpha - 1 - s$.

Having consistent estimators at our disposal, we can propose the following bona fide estimators of the optimal weights,

$$(\hat{w}_1^*, \hat{w}_2^*, \hat{w}_3^*) = \arg \min_{w \in \Delta} w_1^2 \widehat{MSE}(\hat{\Sigma}_n) + w_2^2 \hat{E}_n^\dagger + w_3^2 \hat{E}_n^\diamond + 2w_2 w_3 \hat{D}_n.$$

This yields the feasible shrinkage estimator

$$\hat{\Sigma}_n^{w^*} = w_1^* \hat{\Sigma}_n + w_2^* \hat{\Sigma}_n^\dagger + w_3^* \hat{\Sigma}_n^\diamond.$$

The performance of the corresponding shrinkage estimator $\hat{\Sigma}_n^{w^*}$ is assessed via simulations in the Section 6.

Let us now briefly discuss the case of an interior solution of the minimization problem (5). We have

$$f(w_2, w_3) = (w_2, w_3) Q_n \begin{pmatrix} w_2 \\ w_3 \end{pmatrix} - a_n^\top \begin{pmatrix} w_2 \\ w_3 \end{pmatrix} + \text{MSE}(\hat{\Sigma}_n),$$

where

$$Q_n = \begin{pmatrix} \text{MSE}(\hat{\Sigma}_n) + E^\dagger & \text{MSE}(\hat{\Sigma}_n) + D_n \\ \text{MSE}(\hat{\Sigma}_n) + D_n & \text{MSE}(\hat{\Sigma}_n) + E_n^\diamond \end{pmatrix}, \quad a_n = \begin{pmatrix} \text{MSE}(\hat{\Sigma}_n) \\ \text{MSE}(\hat{\Sigma}_n) \end{pmatrix}.$$

Hence, we are given a quadratic minimization problem under box constraints, and a natural condition is to assume that $\det(Q_n) > 0$. If an interior solution $w^* \in (0, 1)^3$ exists (or the box constraint is omitted), the explicit solution is easily seen to be given by

$$w_2^* = \frac{\text{MSE}(\hat{\Sigma}_n)(E_n^\diamond - D_n)}{F_n}, \quad w_3^* = \frac{\text{MSE}(\hat{\Sigma}_n)(E_n^\dagger - D_n)}{F_n}$$

with

$$F_n = \det(Q_n) = E_n^\dagger E_n^\diamond + \text{MSE}(\hat{\Sigma}_n)(E_n^\dagger + E_n^\diamond - 2D_n) - D_n^2.$$

The resulting optimal weight for the sample covariance matrix $\hat{\Sigma}_n$ is

$$w_1^* = \frac{E_n^\dagger E_n^\diamond - D_n^2}{F_n}.$$

We see from the formula for w_1^* that the optimal solution prefers the nonparametric estimator $\hat{\Sigma}_n$, if the truth Σ_n cannot be well approximated by a banded or Toeplitz matrix. The banded estimator receives a large weight, if the distance between the Toeplitz oracle and Σ_n is large. Analogously, the optimal solution assigns a large weight to the Toeplitz estimator, if Σ_n is not well approximated by the banded oracle, so that $\|\Sigma_n - \Sigma_n^\dagger\|_F^2$ is large. As a result, it is optimal to distribute the weights across the estimators $\hat{\Sigma}_n, \hat{\Sigma}_n^\dagger$ and $\hat{\Sigma}_n^\circ$ according to how well they approximate Σ_n , where the approximation accuracy is measured by the squared distances of the oracles to the truth and, in case of the nonparametric estimator, by the MSE.

Since $D_n^2 \leq E_n^\dagger E_n^\circ$, we have $w_1^* \geq 0$. If $D_n \leq \min(E_n^\dagger, E_n^\circ)$, then $w_2^*, w_3^* \geq 0$ so that $w \in \Delta$. It is interesting to note that $\frac{w_2^*}{w_3^*} = \frac{E_n^\circ - D_n}{E_n^\dagger - D_n}$, i.e., the ratio of the oracle shrinking weights is independent of $\text{MSE}(\hat{\Sigma}_n)$. Further, $w_2^* > w_3^*$ (banding preferable compared to Toeplitz) if and only if

$$E_n^\dagger < E_n^\circ \Leftrightarrow \|\Sigma_n^\dagger - \Sigma_n\|_{F^*} < \|\Sigma_n^\circ - \Sigma_n\|_{F^*}.$$

Even more is true: The above derivations remain formally true without any changes, if we use an arbitrary unbiased estimator $\tilde{\Sigma}_n$ of Σ_n , i.e., an arbitrary measurable function $\tilde{\Sigma}_n$ of the data with $\mathbb{E}(\tilde{\Sigma}_n) = \Sigma_n$. Therefore, we have the following interesting result: The ratio $\frac{w_2^*}{w_3^*}$ is a universal characteristic of the risk minimization problem (4) in the sense that it is independent of the nonparametric estimator and completely determined by Σ_n .

5. Testing for structural breaks

The sequential approximation result of Theorem 3.3 may be applied to test for changes in the covariance matrix $\Sigma_t = \text{Cov}(X_t, X_t) \in \mathbb{R}^{d \times d}$ of a time series X_t . To be precise, we want to test the hypothesis

$$H_0 : \Sigma_t \equiv \Sigma_0 \quad \leftrightarrow \quad H_1 : \Sigma_t \not\equiv \Sigma_0.$$

We propose the following CUSUM-type statistics

$$T_n(w) = \sqrt{n} \max_{k \leq n} \left| v^T \left(\hat{\Sigma}_{k,n}^w - \frac{k}{n} \hat{\Sigma}_{n,n}^w \right) v \right|, \quad w \in \Delta,$$

$$T_n^*(D) = \sup_{w \in D} T_n(w), \quad D \subset \Delta.$$

To determine critical values for the statistic $T_n(w)$ resp. T_n^* , we employ a bootstrap scheme suggested in Mies and Steland (2023). In particular, let $\eta_t \sim \mathcal{N}(0, A_t)$ be independent Gaussian random vectors in \mathbb{R}^3 , where

$$A_t = \frac{1}{b} \left(\sum_{s=t-b+1}^t \chi_s \right)^{\otimes 2}, \quad t = b, \dots, n,$$

$$\chi_s = \left[\nu^j \cdot (X_s X_s^T - \hat{\Sigma}_{n,n}) \right]_{j=1}^3.$$

and $A_t = 0$ for $t < b$. Here, b is a block-length tuning parameter which needs to be chosen suitably, see below. We now define the bootstrapped version of the changepoint statistics as

$$R_n(w) = \sqrt{n} \max_{k \leq n} \left| \frac{1}{n} \sum_{t=b}^k w^T \eta_t - \frac{k}{n} \frac{1}{n} \sum_{t=b}^n w^T \eta_t \right|,$$

$$R_n^*(D) = \sup_{w \in D} R_n(w).$$

Finally, for $\alpha \in (0, 1)$, introduce the conditional quantiles

$$a_n(\alpha, w) = \inf\{a > 0 : P(R_n(w) | X_1, \dots, X_n) \leq \alpha\},$$

$$a_n^*(\alpha, D) = \inf\{a > 0 : P(R_n^*(D) | X_1, \dots, X_n) \leq \alpha\}.$$

We suggest to reject the null hypothesis if

$$T_n(w) > a_n(\alpha, w) + \delta_n, \quad \text{resp.} \quad T_n^*(D) > a_n^*(\alpha, D) + \delta_n,$$

for $\delta_n = 1/(\log n)^p$ for some $p \geq 1$.

Theorem 5.1. *Let $X_t = X_{t,n} = G_t^n(\epsilon_t)$ be an array of d_n -variate time series, such that each kernel G_t^n satisfies (P.1) and (P.2), for some $q > 8$, $\beta > 2$, and with factors $\tilde{\Theta}$ and $\tilde{\Gamma}$ not depending on n . Choose the block length b such that $b = b_n \asymp n^\zeta$ for some $\zeta \in (0, 1)$. If $\text{Cov}(X_{t,n}) = \Sigma_0$ for all $t = 1, \dots, n$, and for some covariance matrix Σ_0 , then*

$$\limsup_{n \rightarrow \infty} P(T_n(w) > a_n(\alpha, w) + \delta_n) \leq \alpha, \quad w \in \Delta,$$

$$\limsup_{n \rightarrow \infty} P(T_n^*(D) > a_n^*(\alpha, D) + \delta_n) \leq \alpha, \quad D \subset \Delta.$$

Hence, the proposed bootstrap scheme indeed maintains the specified size for a change in covariance. A particular feature of this test is that it is robust against nuisance changes, because even under the null hypothesis the time series X_t may be nonstationary, except for its stationary marginal covariance matrix. The relevance of this kind of robustness has been first highlighted by Zhou (2013) for changes in mean, see also Górecki, Horváth and Kokoszka (2018) and Pešta and Wendler (2020). Changes in the second moment structure in the presence of non-constant mean have been studied by Dette, Wu and Zhou (2019) and Schmidt et al. (2021), and a robust CUSUM test for a broad class of parameters is introduced in Mies (2021). Nevertheless, all these references focus on the low-dimensional case, whereas our new test is also applicable in high dimensions.

Despite its robustness, the suggested test still maintains power against classical changepoint alternatives.

Theorem 5.2. *Let $X_t = X_{t,n} = G_t^n(\epsilon_t)$ be an array of d_n -variate time series, such that each kernel G_t^n satisfies (P.1) and (P.2) for some $q > 8$, $\beta > 2$, and with factors $\tilde{\Theta}$ and $\tilde{\Gamma}$ not depending on n . Suppose that for some $u \in (0, 1)$,*

$\Sigma_t = \Sigma_1$ for $t = 1, \dots, \lfloor un \rfloor$, and $\Sigma_t = \Sigma_n$ for larger times $t = \lfloor un \rfloor + 1, \dots, n$. Choose $b = b_n \rightarrow \infty$ such that $b_n \ll n$, and standardize v such that $\|v\|_1 = 1$. Define the change size

$$f_n(w) = \left| \sum_{k=1,2,3} w_k v^k \cdot [\Sigma_n - \Sigma_1] \right|, \quad w \in \Delta,$$

$$f_n(D) = \sup_{w \in D} f_n(w), \quad D \subset \Delta.$$

If $\sqrt{n}f_n(w) \rightarrow \infty$, then, as $n \rightarrow \infty$,

$$P(T_n(w) > a_n(\alpha, w) + \delta_n) \rightarrow 1. \tag{6}$$

If $\sqrt{n}f_n(D) \rightarrow \infty$, and if D has non-empty interior, then, as $n \rightarrow \infty$,

$$P(T_n^*(D) > a_n^*(\alpha, D) + \delta_n) \rightarrow 1. \tag{7}$$

Remark 2. Theorem 5.2 shows that the test detects local alternatives at rate $1/\sqrt{n}$, quantified in a suitable sense. In the course of the proof, we show that $a_n^*(D) = \mathcal{O}_P(1) + o_P(T_n^*(D))$. The latter bound might not be sharp, as the projection via the vector v may lead to some averaging, making $a_n^*(D)$ even smaller. This effect would increase the power of our test, and allow to detect local alternatives at a rate faster than \sqrt{n} , by ‘pooling’ power across the dimensions. However, to make this analysis rigorous, further assumption on the dependence structure of the high-dimensional time series X_t are needed, which we explicitly do not want to impose in this paper.

6. Simulation

To demonstrate the implementation of our proposed changepoint test and the benefits of shrinkage, we assess our methodology for three examples of high-dimensional time series. To emphasize the relevance of accounting for nonstationarity, we purposely simulate weakly stationary time series. However, the autocovariance structure of the outer products $X_t X_t^T$ is non-constant in time, and hence the same holds for the local long run variances Ξ_t in Theorem 3.3.

6.1. Model A

First, we simulate a high-dimensional time series according to the vector ARMA(1,1) model

$$X_t = a\Pi X_{t-1} + b\epsilon_t + \epsilon_{t-1}.$$

The matrix $\Pi \in \mathbb{R}^{d \times d}$ is a permutation matrix, such that the model is non-explosive. Furthermore, for any t , the $(\epsilon_t)_i$, $i = 1, \dots, d$, are chosen as independent, zero mean random variables having a symmetrized Gamma distribution with shape parameter $\alpha(t/n)$, standardized to unit variance. We set

$\alpha(u) = 2 + \sin(2\pi u)$, $u \in [0, 1]$. Thus, the autocovariance structure is stationary, but the time-series is nonstationary. In particular, the long run covariance matrices Ξ_t which occur in Theorem 3.3 are non-constant.

We assess our changepoint procedure under the null hypothesis of no change, and under the alternative where b changes at time $t = \lceil \frac{n}{2} \rceil$. The projection vector v is chosen randomly as $N/\|N\|_1$ for a standard Gaussian random vector $N \sim \mathcal{N}(0, I_{d \times d})$. The threshold for the tapering estimator is chosen as $\tau_n^\dagger = n^{\frac{1}{2\alpha+1}}$, and $\tau_n^\circ = (nd/\log(nd))^{\frac{1}{2\alpha+1}}$. Note that these thresholds are optimal for estimation under the spectral norm, in the class of bandable matrices with decay rate α . In practice, α is unknown, and we use the threshold for $\alpha = 2$ in our simulations. To determine the shrinkage weights, we also use the threshold $\sigma = \sigma_n^\dagger = n^{\frac{1}{2\alpha+1}} \gg \tau_n^\dagger$. For the bootstrap scheme, the additional offset is chosen as $\delta_n = \log(n)^{-4}$, and the block-length is $b_n = \lceil 5n^{0.2} \rceil$.

Table 1 reports the simulated size and power of the bootstrap test for different combinations of n and $d = d_n$, and shrinkage weights $w^{(1)} = (1, 0, 0)$ and $w^{(2)} = (0.3, 0.3, 0.4)$. We also consider the data-driven shrinkage weight w^* as described in Section 4, which is not covered by our Gaussian approximation results. Regarding the choice of the bandwidth for the estimator $\widehat{MSE}(\hat{\Sigma}_n)$, we employ the method of Newey and West (1994) with Bartlett kernel as implemented in the R package SANDWICH (see Zeileis, Köll and Graham (2020)), and we determine a separate bandwidth for each component $\widehat{\text{Var}}(\sqrt{n}(\hat{\Sigma}_n)_{i,j})$.

When simulating the process, we set $a = 0.5$ and $b = 0.5$ (resp. $b = 0.75$ post-change). Based on our simulation results, we find that the test maintains the specified size 10% and is indeed slightly conservative, as established theoretically in Section 5. Moreover, the power of the changepoint test increases with sample size, as expected, but also with increasing dimension. This may be explained by the fact that here, the change affects all coordinates, such that the additional information can be used to improve the detection performance. It is also interesting to observe that the additional usage of the tapered and the Toeplitz estimator further improves the power of the test. This is interesting because the shrinkage estimator has originally been proposed as a method to improve the performance of a point estimate. Our simulation results demonstrate that shrinkage can also be also useful for testing.

6.2. Models B & C

As a second example, we consider the vector ARMA(1,1) process X_t given by

$$X_t = aX_{t-1} + b\epsilon_t + \epsilon_{t-1},$$

with a and b as above, and $\epsilon_t \sim \mathcal{N}(0, A)$ iid random vectors, and we initialize $X_0 \sim \mathcal{N}(0, A(1+b^2)/(1-a^2))$ to ensure weak stationarity. The symmetric positive semidefinite matrix $A \in \mathbb{R}^{d \times d}$ is chosen as $A_{i,j} = \gamma(t_i - t_j)$, for $\gamma(h) = |h+1|^{2H} + |h-1|^{2H} - 2|h|^{2H}$, which is the autocovariance function of fractional Gaussian noise. y For the values t_i , we consider the scenarios (B) $t_i = i$, such that

TABLE 1
 Power (and size), nominal level 10%. p -values based on 10^3 bootstrap samples, reported values based on 10^4 Monte Carlo simulations.

		$n = 100$	$n = 500$	$n = 1000$	$n = 2000$
Model A					
$w^{(1)}$	$d = 4$	0.102 (0.09)	0.220 (0.08)	0.433 (0.09)	0.727 (0.09)
	$d = 4n^{0.3}$	0.101 (0.09)	0.289 (0.08)	0.520 (0.08)	0.824 (0.09)
	$d = 4n^{0.5}$	0.098 (0.07)	0.282 (0.08)	0.535 (0.08)	0.822 (0.09)
$w^{(2)}$	$d = 4$	0.107 (0.08)	0.334 (0.08)	0.594 (0.09)	0.878 (0.10)
	$d = 4n^{0.3}$	0.155 (0.07)	0.823 (0.08)	0.990 (0.09)	1.000 (0.09)
	$d = 4n^{0.5}$	0.201 (0.05)	0.967 (0.07)	1.000 (0.07)	1.000 (0.08)
w^*	$d = 4$	0.102 (0.08)	0.252 (0.08)	0.458 (0.09)	0.751 (0.09)
	$d = 4n^{0.3}$	0.119 (0.08)	0.530 (0.08)	0.826 (0.08)	0.985 (0.09)
	$d = 4n^{0.5}$	0.130 (0.06)	0.697 (0.07)	0.956 (0.07)	1.000 (0.09)
Model B					
$w^{(1)}$	$d = 4$	0.108 (0.08)	0.277 (0.08)	0.503 (0.09)	0.803 (0.09)
	$d = 4n^{0.3}$	0.102 (0.08)	0.275 (0.07)	0.502 (0.08)	0.804 (0.09)
	$d = 4n^{0.5}$	0.101 (0.08)	0.276 (0.07)	0.505 (0.08)	0.798 (0.09)
$w^{(2)}$	$d = 4$	0.137 (0.08)	0.438 (0.08)	0.712 (0.09)	0.938 (0.09)
	$d = 4n^{0.3}$	0.198 (0.08)	0.861 (0.07)	0.993 (0.08)	1.000 (0.09)
	$d = 4n^{0.5}$	0.270 (0.07)	0.966 (0.08)	1.000 (0.08)	1.000 (0.09)
w^*	$d = 4$	0.115 (0.08)	0.302 (0.08)	0.531 (0.09)	0.829 (0.09)
	$d = 4n^{0.3}$	0.133 (0.08)	0.528 (0.07)	0.837 (0.08)	0.989 (0.09)
	$d = 4n^{0.5}$	0.157 (0.07)	0.671 (0.07)	0.941 (0.08)	0.999 (0.09)
Model C					
$w^{(1)}$	$d = 4$	0.098 (0.08)	0.273 (0.08)	0.502 (0.08)	0.806 (0.09)
	$d = 4n^{0.3}$	0.104 (0.08)	0.277 (0.08)	0.510 (0.08)	0.806 (0.09)
	$d = 4n^{0.5}$	0.094 (0.08)	0.269 (0.07)	0.506 (0.08)	0.799 (0.09)
$w^{(2)}$	$d = 4$	0.124 (0.08)	0.415 (0.08)	0.692 (0.09)	0.929 (0.09)
	$d = 4n^{0.3}$	0.170 (0.08)	0.774 (0.08)	0.971 (0.08)	1.000 (0.09)
	$d = 4n^{0.5}$	0.225 (0.07)	0.908 (0.08)	0.997 (0.08)	1.000 (0.09)
w^*	$d = 4$	0.108 (0.08)	0.321 (0.08)	0.573 (0.08)	0.856 (0.09)
	$d = 4n^{0.3}$	0.132 (0.08)	0.460 (0.08)	0.733 (0.08)	0.947 (0.09)
	$d = 4n^{0.5}$	0.143 (0.08)	0.518 (0.07)	0.786 (0.08)	0.952 (0.09)

A is a Toeplitz matrix, and (C) $t_i = \sqrt{i}$, such that A is not a Toeplitz matrix. For the bootstrap scheme and for the determination of shrinkage weights, we use the same settings as in Model A.

The size and power of the changepoint test are presented in Table 1. As for Model (A), the test turns out to be conservative, and gains power with increasing sample size, and higher dimension. In our simulations, we also find the deterministic weight $w^{(2)}$ leading to higher power compared to the data-driven weight w^* . Thus, future work could explore the optimal choice of shrinkage weights from a testing perspective.

Under the null hypothesis of no change, the marginal covariance matrix of X_t is given by $\text{Cov}(X_t) = A \frac{1+2ab+b^2}{1-a^2}$. Thanks to this explicit formula, we are able to analyze the error of estimation of the proposed shrinkage estimator with data-driven shrinkage weights, see Table 2. As benchmarks, we evaluate the performance of the estimators $\hat{\Sigma}_n$, $\hat{\Sigma}_n^\dagger$, and $\hat{\Sigma}_n^\circ$ individually. For model B, the Toeplitz estimator $\hat{\Sigma}_n^\circ$ is found to perform best, which could be expected be-

TABLE 2. Mean square estimation error of the marginal covariance $\Sigma = \text{Cov}(X_t)$ under stationarity, measured in the scaled Frobenius norm $\|\cdot\|_{F^*}$. Errors are reported, in this order, for (i) sample covariance $\hat{\Sigma}_n$, (ii) tapered estimator $\hat{\Sigma}_n^\dagger$, (iii) Toeplitz-type estimator $\hat{\Sigma}_n^\diamond$, (iv) the shrinkage estimator $\hat{\Sigma}_n^{w^*}$. Reported values based on 10^4 Monte Carlo simulations.

	$n = 100$	$n = 500$	$n = 1000$	$n = 2000$
Model B				
$d = 4$	2.50 2.32 0.87 2.34	0.52 0.48 0.18 0.48	0.26 0.24 0.09 0.24	0.13 0.12 0.05 0.12
$d = 4n^{0.3}$	6.38 3.84 0.78 4.01	2.55 1.29 0.16 1.42	1.49 0.76 0.08 0.82	0.95 0.49 0.04 0.53
$d = 4n^{0.5}$	19.96 4.32 0.60 7.71	9.06 1.43 0.15 3.17	6.39 0.86 0.08 2.13	4.54 0.56 0.04 1.48
Model C				
$d = 4$	2.63 2.61 1.49 2.45	0.55 0.67 0.61 0.52	0.28 0.42 0.50 0.27	0.14 0.29 0.44 0.14
$d = 4n^{0.3}$	6.85 5.17 3.91 4.80	2.73 2.55 5.21 1.98	1.58 1.79 5.69 1.23	1.01 1.26 6.80 0.82
$d = 4n^{0.5}$	21.25 12.47 9.05 10.64	9.48 11.13 12.98 5.75	6.66 12.37 15.69 4.55	4.70 12.27 19.12 3.53

cause the true covariance matrix is of Toeplitz type. On the other hand, the shrinkage estimator always improves upon the sample covariance matrix $\hat{\Sigma}_n$, and is often close the performance of the tapered estimator $\hat{\Sigma}_n^\dagger$. For model C, the decay of the off-diagonal entries of the true covariance matrix is slower, hence the performance of $\hat{\Sigma}_n^\dagger$ is also worse. Indeed, its estimation error is larger than the error of the sample covariance $\hat{\Sigma}_n$. Yet, for $n \geq 500$, the error of the shrinkage estimator is lower than that of all other estimators. This demonstrates that our proposed shrinkage estimator not only chooses among the three estimators, but may indeed improve the performance even further due to the convex combination.

7. Proofs

Proof of Lemma 2.1. The result follows from Jensen’s inequality:

$$\begin{aligned} \mathbb{E} \left| \sum_{j=1}^d v_j X_{tj} \right|^q &\leq \|v\|_1^q \mathbb{E} \left(\sum_{j=1}^d \frac{|v_j|}{\|v\|_1} |X_{tj}| \right)^q \\ &\leq \|v\|_1^q \sum_{j=1}^d \frac{|v_j|}{\|v\|_1} \mathbb{E}|X_{tj}|^q = \|v\|_1^q \max_{1 \leq j \leq d} \mathbb{E}|X_{tj}|^q \end{aligned}$$

□

Proof of Theorem 2.2. The time series Z_t may be written as $Z_t = \tilde{G}_t(\epsilon_t)$, with kernel $\tilde{G}_t = VG_t$, such that we may apply Theorem 3.1 of [Mies and Steland \(2023\)](#). It can be verified that \tilde{G}_t satisfies conditions (G.1) and (G.2) therein, with $\Gamma = \tilde{\Gamma}$ and $\Theta = \sqrt{m}\tilde{\Theta}\|V\|_\infty$. To see this, note that for any m -variate random variable X , we have

$$(\mathbb{E}\|X\|_2^q)^{\frac{1}{q}} \leq (\mathbb{E}\|X\|_q^q)^{\frac{1}{q}} m^{\frac{1}{2} - \frac{1}{q}} \leq m^{\frac{1}{2}} \max_{l=1, \dots, m} (\mathbb{E}|X_l|^q)^{\frac{1}{q}},$$

because $\|x\|_2 \leq m^{\frac{1}{2} - \frac{1}{q}} \|x\|_q$ for any vector $x \in \mathbb{R}^m$. Moreover, for any $l = 1, \dots, m$,

$$\begin{aligned} (\mathbb{E}|V_{l,\cdot}G_t(\epsilon_0)|^q)^{\frac{1}{q}} &= \left(\mathbb{E} \left| \sum_{r=1}^d V_{l,r}G_{t,r}(\epsilon_0) \right|^q \right)^{\frac{1}{q}} \\ &\leq \sum_{r=1}^d |V_{l,r}| (\mathbb{E}|G_{t,r}(\epsilon_0)|^q)^{\frac{1}{q}} \leq \|V\|_\infty \tilde{\Theta}, \end{aligned}$$

such that

$$(\mathbb{E}\|VG_t(\epsilon_0)\|^q)^{\frac{1}{q}} \leq \sqrt{m}\|V\|_\infty \tilde{\Theta}.$$

The same argument can be used to establish (G.1) and (G.2). □

Proof of Proposition 3.1. Observe that for four random variables $Y_1, Y_2, Z_1, Z_2 \in L_{2q}$, it holds that

$$\begin{aligned} (\mathbb{E}|Y_1 Z_1 - Y_2 Z_2|^q)^{\frac{1}{q}} &\leq (\mathbb{E}|(Y_1 - Y_2) Z_1|^q)^{\frac{1}{q}} + (\mathbb{E}|Y_2 (Z_1 - Z_2)|^q)^{\frac{1}{q}} \\ &\leq (\mathbb{E}|Y_1 - Y_2|^{2q})^{\frac{1}{2q}} (\mathbb{E}|Z_1|^{2q})^{\frac{1}{2q}} + (\mathbb{E}|Y_2|^{2q})^{\frac{1}{2q}} (\mathbb{E}|Z_1 - Z_2|^{2q})^{\frac{1}{2q}}. \end{aligned}$$

Inequality (P.1) may be obtained by setting $Y_1 = G_{t,l}(\epsilon_t), Y_2 = G_{t,l}(\tilde{\epsilon}_{t,t-j})$ and $Z_1 = G_{t,l'}(\epsilon_t), Z_2 = G_{t,l'}(\tilde{\epsilon}_{t,t-j})$. Inequality (P.2) may be obtained by setting $Y_1 = G_{t,l}(\epsilon_t), Y_2 = G_{t-1,l}(\epsilon_t)$ and $Z_1 = G_{t,l'}(\epsilon_t), Z_2 = G_{t-1,l'}(\epsilon_t)$. \square

Proof of Theorem 3.2. By (Mies and Steland, 2023, Th. 3.2) we have under Assumption (P.1) for all d, n

$$\left(\mathbb{E} \max_{k \leq n} \left| \sum_{t=1}^k X_{ti} X_{tj} - (\text{Var}(X_t))_{ij} \right|^2 \right)^{\frac{1}{2}} \leq C \frac{\tilde{\Theta}^2}{\beta - 1} n^{\frac{1}{2}}, \quad 1 \leq i, j \leq d,$$

such that $\hat{\Sigma}_{n,ij} = (\hat{\Sigma}_n)_{ij}$ satisfies

$$\max_{1 \leq i, j \leq d} \mathbb{E}(\hat{\Sigma}_{n,ij} - \Sigma_{n,ij})^2 \leq C \frac{\tilde{\Theta}^4}{(\beta - 1)^2} n^{-1}. \tag{8}$$

To show the bound for $\hat{\Sigma}_n^\dagger$, observe that

$$\mathbb{E} \|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2 \leq 4(A_n + B_n + C_n)$$

where

$$\begin{aligned} A_n &= d^{-1} \sum_{m \leq \tau/2} \sum_{|i-j|=m} \mathbb{E}(\hat{\Sigma}_{n,ij} - \Sigma_{n,ij})^2, \\ B_n &= d^{-1} \sum_{\tau/2 < m \leq \tau} \sum_{|i-j|=m} \omega(|i-j|)^2 \mathbb{E}(\hat{\Sigma}_{n,ij} - \Sigma_{n,ij})^2, \\ C_n &= d^{-1} \sum_{\tau/2 < m \leq \tau} \sum_{|i-j|=m} (1 - \omega(|i-j|))^2 \Sigma_{n,ij}^2, \\ D_n &= d^{-1} \sum_{m > \tau} \sum_{|i-j|=m} \Sigma_{n,ij}^2. \end{aligned}$$

Clearly, $|C_n + D_n| = O(\tau^{-2\alpha-1})$. Furthermore,

$$|A_n + B_n| \leq C \frac{\tilde{\Theta}^4}{(\beta - 1)^2} \frac{1}{nd} \sum_{m \leq \tau} \sum_{|i-j|=m} 1 \leq \frac{C}{n} \min(\tau, d).$$

Thus,

$$\mathbb{E} \|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2 \leq C \left\{ \tau^{-2\alpha-1} + \frac{\tilde{\Theta}^4}{(\beta - 1)^2} \frac{\min(\tau, d)}{n} \right\}.$$

The optimal rate is attained for $\tau = \min \left[(n(\beta - 1)^2 / \tilde{\Theta}^4)^{\frac{1}{2\alpha+2}}, d \right]$.

To show the bound for $\hat{\Sigma}_n^\circ$ observe that by Jensen's inequality

$$\begin{aligned} \mathbb{E}(\hat{\sigma}_m - \sigma_m)^2 &= \mathbb{E}\left(\frac{1}{d-m} \sum_{k-l=m} \hat{\Sigma}_{n,kl} - \Sigma_{n,kl}\right)^2 \\ &\leq \frac{1}{d-m} \sum_{k-l=m} \mathbb{E}(\hat{\Sigma}_{n,kl} - \Sigma_{n,kl})^2 \\ &\leq C \frac{\tilde{\Theta}^4}{(\beta-1)^2} n^{-1}, \end{aligned}$$

where $\sigma_m = \frac{1}{d-m} \sum_{k-l=m} \Sigma_{n,kl}$ for all i, j with $i - j = m \leq \tau$. We have

$$\begin{aligned} \mathbb{E}\|\hat{\Sigma}_n^\circ - \Sigma_n\|_{F,*}^2 &\leq \frac{2}{d} \sum_{m \leq \tau} \sum_{i-j=m} \mathbb{E}(\hat{\Sigma}_{n,ij}^\circ - \Sigma_{n,ij})^2 + \frac{2}{d} \sum_{m > \tau} \sum_{i-j=m} (\Sigma_{n,ij})^2 \\ &= \frac{2}{d} \sum_{m \leq \tau} \sum_{i-j=m} \mathbb{E}(\omega(i-j)(\hat{\sigma}_m - \sigma_m))^2 + \frac{2}{d} \sum_{m > \tau} \sum_{i-j=m} (\Sigma_{n,ij})^2 \\ &= G_n + H_n. \end{aligned}$$

Since $\Sigma_{n,ij}^\circ \leq C m^{-\alpha-1}$, we have $\sigma_m \leq C m^{-\alpha-1}$ for $i - j = m$. Thus, $H_n = O(\tau^{-2\alpha-1})$. Moreover, $G_n = O(\min(d, \tau)/n)$. We obtain the same bound as in case (i), i.e.

$$\mathbb{E}\|\hat{\Sigma}_n^\circ - \Sigma_n\|_{F,*}^2 \leq C \left\{ \tau^{-2\alpha-1} + \frac{\tilde{\Theta}^4}{(\beta-1)^2} \frac{\min(\tau, d)}{n} \right\}.$$

Optimal choice of τ yields the same rates as in case (i). □

Proof of Lemma 4.1. One only needs to verify Condition 2 of [Sancetta \(2008\)](#). Essentially, this is specific Doukhan-Louhichi type weak dependence condition. It is known that Bernoulli shifts satisfy such weak dependence conditions, but the applying the results in the literature would require to impose stronger assumptions on G_t . Therefore, we give a direct proof working under (P.1). Let $\bar{\epsilon}_{i,j} = (\epsilon_i, \dots, \epsilon_{j+1}, \epsilon'_j, \epsilon'_{j-1}, \dots)$ where $\{\epsilon'_t\}$ is an independent copy of $\{\epsilon_t\}$. Fix $u, v \in \{1, \dots, 4\}$, $(i_1, \dots, i_v), (s_1, \dots, s_v) \in \mathbb{N}^v$ and $(j_1, \dots, j_u), (t_1, \dots, t_u) \in \mathbb{N}^u$ such that $s_1 \leq \dots \leq s_u + r \leq t_1 \leq \dots \leq t_v$ for some $r \in \mathbb{N}$. Put $U_{t_v}(\epsilon_{t_v}) = X_{t_1, j_1} \dots X_{t_v, j_v}$ and $V_{s_u}(\epsilon_{s_u}) = Y_{s_1, i_1} \dots Y_{s_u, i_u}$. Then

$$\begin{aligned} \text{Cov}(U_{t_v}(\epsilon_{t_v}), V_{s_u}(\epsilon_{s_u})) &= \text{Cov}(U_{t_v}(\epsilon_{t_v-s_u}), V_{s_u}(\epsilon_0)) \\ &= \text{Cov}(U_{t_v}(\bar{\epsilon}_{t_v-s_u,0}), V_{s_u}(\epsilon_0)) \\ &\quad + \text{Cov}(U_{t_v}(\epsilon_{s_u,0}) - U_{t_v}(\bar{\epsilon}_{t_v-s_u,0}), V_{s_u}(\epsilon_0)) \\ &= \text{Cov}(U_{t_v}(\epsilon_{s_u,0}) - U_{t_v}(\bar{\epsilon}_{t_v-s_u,0}), V_{s_u}(\epsilon_0)) \end{aligned}$$

By telescoping over the random variables $U_{t_v}(l) = U_{t_v}(\epsilon_{t_v-s_u}, \dots, \epsilon_{t_v-s_u-l+1}, \bar{\epsilon}_{t_v-s_u-l})$, so that $\sum_{l > t_v-s_u} U_{t_v}(l) - U_{t_v}(l-1) = U_{t_v}(\epsilon_{s_u,0}) - U_{t_v}(\bar{\epsilon}_{t_v-s_u,0})$, and

noting that $\mathbb{E}|(U_{t_v}(l) - U_{t_v}(l-1))V_{s_u}(\epsilon_0)| \leq \Theta l^{-\beta}$ for some constant Θ' by (P.1), we obtain

$$\text{Cov}(U_{t_v}(\epsilon_{t_v}), V_{s_u}(\epsilon_{s_u})) = \mathcal{O}(r^{-\beta+1}),$$

which establishes the Doukhan-Louhichi type weak dependence condition required in Condition 2 of Sancetta (2008), since $\beta > 2$. \square

Proof of Theorem 4.2. The first assertion follows from rearranging terms,

$$\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger - (\Sigma_n^\circ - \Sigma_n) = (\hat{\Sigma}_n^\circ - \Sigma_n^\circ) - (\hat{\Sigma}_n^\dagger - \Sigma_n)$$

and the triangle inequality, so that by Theorem 3.2

$$\begin{aligned} \mathbb{E}\|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger - (\Sigma_n^\circ - \Sigma_n)\|_{F^*} &\leq \mathbb{E}\|\hat{\Sigma}_n^\circ - \Sigma_n^\circ\|_{F^*} + \mathbb{E}\|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*} \\ &\leq \left(\mathbb{E}\|\hat{\Sigma}_n^\circ - \Sigma_n^\circ\|_{F^*}^2\right)^{\frac{1}{2}} + \left(\mathbb{E}\|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2\right)^{\frac{1}{2}} \\ &= \mathcal{O}\left(n^{-\frac{\alpha+1/2}{2\alpha+2}}\right), \end{aligned}$$

and, by the c_r -inequality,

$$\mathbb{E}\|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger - (\Sigma_n^\circ - \Sigma_n)\|_{F^*}^2 \leq 2\mathbb{E}\|\hat{\Sigma}_n^\circ - \Sigma_n^\circ\|_{F^*}^2 + 2\mathbb{E}\|\hat{\Sigma}_n^\dagger - \Sigma_n\|_{F^*}^2 = \mathcal{O}\left(n^{-\frac{2\alpha+1}{2\alpha+2}}\right).$$

For the second assertion of (i), use the fact that if $A_n - B_n \leq C_n$ on $A_n \geq B_n$ and $B_n - A_n \leq D_n$ on $A_n < B_n$, for random variables A_n, B_n and bounds C_n, D_n , then

$$\begin{aligned} \mathbb{E}|A_n - B_n| &\leq \mathbb{E}(A_n - B_n)1_{A_n \geq B_n} + \mathbb{E}(-A_n + B_n)1_{A_n < B_n} \\ &\leq \mathbb{E}C_n + \mathbb{E}D_n \end{aligned}$$

Apply this bound with $A_n = \|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger\|_{F^*}$, $B_n = \|\Sigma_n^\circ - \Sigma_n\|_{F^*}$ and the bounds

$$\|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger\|_{F^*} - \|\Sigma_n^\circ - \Sigma_n\|_{F^*} \leq \|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger - (\Sigma_n^\circ - \Sigma_n)\|_{F^*} = C_n$$

as well as

$$\|\Sigma_n^\circ - \Sigma_n\|_{F^*} - \|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger\|_{F^*} \leq \|\Sigma_n^\circ - \Sigma_n - (\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger)\|_{F^*} = D_n,$$

which follow from the triangle inequality. Therefore,

$$\mathbb{E}\left|\hat{E}_n^\circ - E_n^\circ\right| = \mathbb{E}\left|\|\hat{\Sigma}_n^\circ - \hat{\Sigma}_n^\dagger\|_{F^*} - \|\Sigma_n^\circ - \Sigma_n\|_{F^*}\right| = \mathcal{O}\left(n^{-\frac{\alpha+1/2}{2\alpha+2}}\right).$$

To show (ii) observe that we have, similarly as above, the bounds

$$\pm(\|\hat{\Delta}_n^\dagger\|_{F^*} - \|\Delta_n^\dagger\|_{F^*}) \leq \|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}$$

yielding

$$\mathbb{E}\left|\|\hat{\Delta}_n^\dagger\|_{F^*} - \|\Delta_n^\dagger\|_{F^*}\right| \leq 2\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*} \leq 2\sqrt{\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2}.$$

Thus, it suffices to study $\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2$. We have

$$\begin{aligned} \mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2 &= d^{-1} \sum_{\tau^\dagger/2 \leq m \leq \tau^\dagger} \sum_{|i-j|=m} (1 - \omega(|i-j|))^2 \mathbb{E}(\hat{\Sigma}_{n,ij} - \Sigma_{n,ij})^2 \\ &\quad + d^{-1} \sum_{\tau^\dagger < m \leq \sigma^\dagger} \sum_{|i-j|=m} \mathbb{E}(\hat{\Sigma}_{n,ij} - \Sigma_{n,ij})^2 \\ &\quad + d^{-1} \sum_{m > \sigma^\dagger} \sum_{|i-j|=m} \Sigma_{n,ij}^2 \\ &= \mathcal{O}(\sigma^\dagger n^{-1}) + \mathcal{O}((\sigma^\dagger)^{-2\alpha-1}) \\ &= \mathcal{O}\left(n^{-\frac{2\alpha+1-s}{2\alpha+2}}\right) + \mathcal{O}\left(n^{-\frac{(s+1)(2\alpha+1)}{2\alpha+2}}\right). \end{aligned}$$

Both terms are $o(1)$ and the first term dominates because the choice of σ^\dagger is essentially undersmoothing. Therefore,

$$\mathbb{E}\left|\|\hat{\Delta}_n^\dagger\|_{F^*} - \|\Delta_n^\dagger\|_{F^*}\right| = \mathcal{O}(\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}) = \mathcal{O}\left(\sqrt{\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2}\right) = \mathcal{O}(r_n),$$

where $r_n = n^{-\frac{2\alpha+1-s}{2(2\alpha+2)}}$.

(iii) First note that

$$\|\Delta_n^\dagger\|_{F^*}^2 = \|\Sigma_n - \Sigma_n^\dagger\|_{F^*}^2 \leq \frac{1}{d} \sum_{|i-j| > \frac{\tau}{2}} \Sigma_{n,ij}^2 = \mathcal{O}((\tau_n^\dagger)^{-2\alpha-1}) = \mathcal{O}\left(n^{-\frac{3\alpha+1}{2\alpha+2}}\right)$$

and

$$\begin{aligned} \|\Delta_n^\diamond\|_{F^*}^2 &= \|\Sigma_n - \Sigma_n^\diamond\|_{F^*}^2 = \frac{1}{d} \sum_{|i-j| \leq \tau} \omega(|i-j|)^2 (\Sigma_{n,ij} - \sigma_{|i-j|})^2 + \frac{1}{d} \sum_{|i-j| > \tau} \Sigma_{n,ij}^2 \\ &= \mathcal{O}\left(\frac{\tau_n^\diamond(d - \tau_n^\diamond)}{d}\right) + \mathcal{O}((\tau_n^\diamond)^{-2\alpha-1}) \\ &= \mathcal{O}(\tau_n^\diamond). \end{aligned}$$

Using the decomposition

$$\langle \hat{\Delta}_n^\dagger, \hat{\Delta}_n^\diamond \rangle_* - \langle \Delta_n^\dagger, \Delta_n^\diamond \rangle_* = \langle \hat{\Delta}_n^\dagger - \Delta_n^\dagger, \hat{\Delta}_n^\diamond \rangle_* + \langle \Delta_n^\dagger, \hat{\Delta}_n^\diamond - \Delta_n^\diamond \rangle_*$$

we obtain the bound

$$\begin{aligned} \left| \langle \hat{\Delta}_n^\dagger, \hat{\Delta}_n^\diamond \rangle_* - \langle \Delta_n^\dagger, \Delta_n^\diamond \rangle_* \right| &\leq \|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*} (\|\Delta_n^\diamond\|_{F^*} + \|\hat{\Delta}_n^\diamond - \Delta_n^\diamond\|_{F^*}) \\ &\quad + \|\Delta_n^\dagger\|_{F^*} \|\hat{\Delta}_n^\diamond - \Delta_n^\diamond\|_{F^*} \end{aligned}$$

leading to

$$\begin{aligned} \mathbb{E}\left| \langle \hat{\Delta}_n^\dagger, \hat{\Delta}_n^\diamond \rangle_* - \langle \Delta_n^\dagger, \Delta_n^\diamond \rangle_* \right| &\leq \|\Delta_n^\diamond\|_{F^*} \mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*} \\ &\quad + \sqrt{\mathbb{E}\|\hat{\Delta}_n^\dagger - \Delta_n^\dagger\|_{F^*}^2} \sqrt{\mathbb{E}\|\hat{\Delta}_n^\diamond - \Delta_n^\diamond\|_{F^*}^2} \\ &\quad + \|\Delta_n^\dagger\|_{F^*} \mathbb{E}\|\hat{\Delta}_n^\diamond - \Delta_n^\diamond\|_{F^*} \end{aligned}$$

$$\begin{aligned}
 &= \mathcal{O}(\sqrt{\tau_n^\circ} r_n) + \mathcal{O}(r_n n^{-\frac{\alpha+1/2}{2\alpha+2}}) + \mathcal{O}((\tau_n^\dagger)^{-2\alpha-1} r_n) \\
 &= \mathcal{O}(\sqrt{\tau_n^\circ} r_n) \\
 &= \mathcal{O}\left(\sqrt{\left(\frac{d}{\log(nd)}\right)^{\frac{1}{2\alpha+1}} n^{\frac{1}{2\alpha+1} + \frac{s+1}{2\alpha+2} - 1}}\right) \\
 &\leq \mathcal{O}\left(\sqrt{\left(\frac{d}{\log(nd)}\right)^{\frac{1}{2\alpha+1}} n^{\frac{s+2}{2\alpha+1} - 1}}\right) \\
 &= \mathcal{O}\left(\sqrt{\left(\frac{d}{\log(nd)}\right)^{\frac{1}{2\alpha+1}} n^{\frac{s+1-2\alpha}{2\alpha+1}}}\right).
 \end{aligned}$$

The latter estimate can be bounded by $\mathcal{O}(d^{\frac{1}{2\alpha+1}} n^{-\frac{2\alpha-s}{2\alpha+2}})$. Hence, if $d = \mathcal{O}(n^\kappa)$ for some $\kappa > 0$, the above bound is $o(1)$ provided $\kappa \leq 2\alpha - 1 - s$. This completes the proof. \square

Proof of Theorem 5.1. Note that the statistics T_n and the critical values $a_n(\alpha, w)$ are both quadratic in $\|v\|_1$, such that we may suppose without loss of generality that $\|v\|_1 = 1$. If we replace the χ_s by the centered random vectors $\bar{\chi}_s = [\nu^j \cdot (X_s X_s^T - \Sigma_0)]_{j=1}^3$, then Theorem 5.1 is a special case of (Mies and Steland, 2023, Prop. 5.3). We briefly show that the estimation error of $\hat{\Sigma}_n$ versus the true covariance matrix is negligible. To this end, denote the sequential estimators of the asymptotic covariance, for $k = b, \dots, n$, by

$$\hat{Q}(k) = \sum_{t=b}^k \frac{1}{b} \left(\sum_{s=t-b+1}^t \chi_s \right)^{\otimes 2}, \quad \bar{Q}(k) = \sum_{t=b}^k \frac{1}{b} \left(\sum_{s=t-b+1}^t \bar{\chi}_s \right)^{\otimes 2}.$$

It suffices to show that

$$\mathbb{E} \max_{k=1, \dots, n} \|\hat{Q}(k) - \bar{Q}(k)\|_{\text{tr}} = \mathcal{O}(\sqrt{nb}), \tag{9}$$

where $\|A\|_{\text{tr}} = \text{tr}((AA^T)^{\frac{1}{2}})$ denotes the trace norm of a symmetric matrix. Since the matrices $\hat{Q}(k)$ and $\bar{Q}(k)$ are always of fixed dimension 3×3 , the choice of norm is in fact irrelevant as all norms are equivalent, and we may use an arbitrary matrix norm $\|\cdot\|$. If (9) holds, then both $\hat{Q}(k)$ and $\bar{Q}(k)$ satisfy Theorem 5.1 in Mies and Steland (2023), such that the bootstrap scheme based on $\hat{Q}(k)$ is valid.

To establish (9), denote $\Delta_n = [\nu^j \cdot (\Sigma_0 - \hat{\Sigma}_{n,n})]_{j=1}^3 = \chi_s - \bar{\chi}_s \in \mathbb{R}^3$, for all s . Via quadratic expansion, we find that, for some universal C which may change from line to line,

$$\begin{aligned}
 &\mathbb{E} \max_k \left\| \hat{Q}(k) - \bar{Q}(k) \right\| \\
 &= \mathbb{E} \max_k \left\| \sum_{t=b}^k \frac{1}{b} \left[\left(\sum_{s=t-b+1}^t (\bar{\chi}_s + \Delta_n) \right)^{\otimes 2} - \left(\sum_{s=t-b+1}^t \bar{\chi}_s \right)^{\otimes 2} \right] \right\|
 \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E} \max_k \left\| \sum_{t=b}^k \frac{1}{b} \left[\left(\sum_{s=t-b+1}^t (\bar{\chi}_s (b\Delta_n)^T + (b\Delta_n) \bar{\chi}_s^T + (b\Delta_n)^{\otimes 2}) \right) \right] \right\| \\
 &\leq C \mathbb{E} \sum_{t=b}^n \left[2 \left\| \sum_{s=t-b+1}^t \bar{\chi}_s \Delta_n^T \right\| + b \|\Delta_n\|^2 \right] \\
 &\leq C \sum_{t=b}^n \sqrt{\mathbb{E} \left\| \sum_{s=t-b+1}^t \bar{\chi}_s \right\|^2} \sqrt{\mathbb{E} \|\Delta_n\|^2} + nb \mathbb{E} \|\Delta_n\|^2.
 \end{aligned}$$

By virtue of Proposition 3.1 and Lemma 2.1, the centered time series $\bar{\chi}_s$ satisfies condition (G.1) of Mies and Steland (2023), such that the Rosenthal-type inequality, Theorem 3.2 therein, is applicable. This yields,

$$\sqrt{\mathbb{E} \left\| \sum_{s=t-b+1}^t \bar{\chi}_s \right\|^2} \leq C\sqrt{b} \sum_{j=1}^{\infty} \tilde{\Theta} j^{-\beta} = \mathcal{O}(\sqrt{b}),$$

because $\tilde{\Theta}$ is fixed in the asymptotic regime of Theorem 5.1. Analogously,

$$\sqrt{\mathbb{E} \|\Delta_n\|^2} = \sqrt{\mathbb{E} \left\| \frac{1}{n} \sum_{t=1}^n \bar{\chi}_t \right\|^2} = \mathcal{O}(1/\sqrt{n}).$$

Hence, $\mathbb{E} \max_k \left\| \hat{Q}(k) - \bar{Q}(k) \right\| = \mathcal{O}(\sqrt{nb} + b) = \mathcal{O}(\sqrt{nb})$. The last inequality holds because $b \leq n$, and establishes (9). \square

Proof of Theorem 5.2. Denote

$$\Sigma_{k,n} := \mathbb{E} \hat{\Sigma}_{k,n} = \frac{k \wedge \lfloor un \rfloor}{n} \Sigma_1 + \frac{(k - \lfloor un \rfloor) \vee 0}{n} \Sigma_n.$$

Theorem 3.3 yields that

$$T_n(w) = \sqrt{n} \max_{k \leq n} \left| \sum_{k=1,2,3} w_k \nu^k \cdot [\Sigma_{k,n} - \Sigma_{n,n}] \right| + \mathcal{O} \left(\sqrt{\frac{1}{n} \sum_{t=1}^n \|\Xi_t\|^2} \right) + o(1),$$

uniformly in $w \in \Delta$. The matrices $\Xi_t \in \mathbb{R}^{3 \times 3}$ are bounded under our assumptions, because (P.1) implies that

$$\text{Cov} (G_{t,j}^n(\epsilon_0) G_{t,i}^n(\epsilon_0), G_{t,j}^n(\epsilon_h) G_{t,i}^n(\epsilon_h)) \leq C \tilde{\Theta}^2 |h|^{1-\beta}, \quad i, j = 1, \dots, d_n.$$

As $\|\nu^k\|_1 \leq 1$ is bounded, we find that $\|\Xi_t\| = \mathcal{O}(1)$. Hence, for some universal $c = c(u) > 0$ depending on the change location $u \in (0, 1)$,

$$T_n(w) \geq c \cdot \sqrt{n} f_n(w) + \mathcal{O}(1).$$

This approximation holds uniformly in $w \in \Delta$, and thus

$$T_n^*(D) \geq \tilde{c} \cdot \sqrt{n} f_n(D) + \mathcal{O}(1).$$

As a second step, we derive an upper bound on the bootstrap quantile $a_n(\alpha, w)$ $a_n^*(\alpha, D)$. Note that, conditionally on X_1, \dots, X_n , the η_t are independent centered random vectors. Hence, the Burkholder-Davis-Gundy inequality yields

$$\mathbb{E}^* |R_n(w)|^2 \leq 4n \mathbb{E}^* \max_{k=1, \dots, n} \left| \frac{1}{n} \sum_{t=1}^k w^T \eta_t \right|^2 \tag{10}$$

$$\leq \frac{C}{n} \sum_{t=b}^n \mathbb{E}^* |w^T \eta_t|^2 \leq \tilde{C} \left\| \frac{1}{n} \sum_{t=b}^n w^T A_t w \right\|, \tag{11}$$

where \mathbb{E}^* denotes the conditional expectation given X_1, \dots, X_n , for brevity.

Let $\bar{\chi}_s = [\nu^j \cdot (X_s X_s^T - \mathbb{E}(X_s X_s^T))]_{j=1}^3$, and define $\bar{A}_t = \frac{1}{b} \left(\sum_{s=t-b+1}^t \bar{\chi}_s \right)^{\otimes 2}$. By virtue of Proposition 3.1, and since $\|\nu^j\|_1$ is bounded, the sequence $\bar{\chi}_t$ satisfies the conditions of (Mies and Steland, 2023, Thm. 5.1), with $d = 3$, which yields that

$$\left\| \frac{1}{n} \sum_{t=b}^n \bar{A}_t - \frac{1}{n} \sum_{t=1}^n \Xi_t \right\| = \mathcal{O}_P \left(\sqrt{\frac{b}{n}} + b^{-1} + b^{2-\beta} \right) = o_P(1).$$

Since Ξ_t is bounded, we find that

$$\frac{1}{n} \sum_{t=b}^n \bar{A}_t = \mathcal{O}_P(1). \tag{12}$$

Now denote $\Delta_{t,n} = \chi_t - \bar{\chi}_t = [\nu^j \cdot (\mathbb{E}(X_s X_s^T) - \hat{\Sigma}_{n,n})]_{j=1}^3$. Similar to the proof of Theorem 5.1, we derive

$$\begin{aligned} \mathbb{E} \left| \frac{1}{n} \sum_{t=b}^n w^T [A_t - \bar{A}_t] w \right| &= \frac{1}{n} \mathbb{E} \left| \sum_{t=b}^n \frac{1}{b} w^T \left[\left(\sum_{s=t-b+1}^t (\bar{\chi}_s + \Delta_{s,n}) \right)^{\otimes 2} \right. \right. \\ &\quad \left. \left. - \left(\sum_{s=t-b+1}^t \bar{\chi}_s \right)^{\otimes 2} \right] w \right| \\ &\leq \frac{2}{n} \mathbb{E} \left| \sum_{t=b}^n \frac{1}{b} w^T \left(\sum_{s=t-b+1}^t \bar{\chi}_s \right) \left(\sum_{s=t-b+1}^t \Delta_{s,n} \right)^T w \right| \\ &\quad + \frac{1}{n} \mathbb{E} \left| \sum_{t=b}^n \frac{1}{b} w^T \left(\sum_{s=t-b+1}^t \Delta_{s,n} \right)^{\otimes 2} w \right| \end{aligned}$$

$$\begin{aligned} &\leq \frac{2}{n} \sum_{t=b}^n \mathbb{E} \left[\left\| \sum_{s=t-b+1}^t \bar{\chi}_s \right\| \cdot |\Delta_n^*(w)| \right] + b\mathbb{E}|\Delta_n^*(w)|^2 \\ &\leq 2\sqrt{b\mathbb{E}|\Delta_n^*(w)|^2} \max_{t=b, \dots, n} \sqrt{\mathbb{E} \left\| \frac{1}{\sqrt{b}} \sum_{s=t-b+1}^t \bar{\chi}_s \right\|^2} \\ &\quad + b\mathbb{E}|\Delta_n^*(w)|^2, \end{aligned}$$

for $\Delta_n^*(w) = \max_{s=1, \dots, n} |w^T \Delta_{s,n}|$. By virtue of (P.1), we may conclude that the autocovariances $\|\text{Cov}(\chi_r, \chi_s)\| \leq C\tilde{\Theta}^2|r-s|^{1-\beta}$ are summable (Mies and Steland, 2023, Prop. 6.4), and we hence conclude that

$$\max_{t=b, \dots, n} \sqrt{\mathbb{E} \left\| \frac{1}{\sqrt{b}} \sum_{s=t-b+1}^t \bar{\chi}_s \right\|^2} = \mathcal{O}(1).$$

The same bound on the autocovariances yields $\text{Var}(\hat{\Sigma}_{n,n}) = \mathcal{O}(1/n)$. Thus, via the classical bias-variance-decomposition, $\mathbb{E}|\Delta_n^*(w)|^2 \leq C(\frac{1}{n} + f_n(w)^2)$, uniformly in $w \in \Delta$. We conclude that

$$\left\| \frac{1}{n} \sum_{t=b}^n w^T [A_t - \bar{A}_t] w \right\| = \mathcal{O}_P \left(\sqrt{b} f_n(w) \right) + o_P(1).$$

Together with (11) and (12), we obtain $R_n(w) = \mathcal{O}_P \left(\sqrt{b} f_n(w) \right) + o_P(1)$, and hence $a_n(\alpha, w) = \mathcal{O}_P \left(\sqrt{b} f_n(w) \right) + o_P(1)$. In combination with the lower bound on $T_n(w)$, and $b = b_n \ll n$, this establishes (6)

To proof (7), we need an upper bound on the supremum $R_n(D) = \sup_{w \in D} R_n(w)$. Since $w \in \Delta$ may be regarded as convex weights, the triangle inequality yields

$$R_n(D) \leq R_n(w_1) + R_n(w_2) + R_n(w_3) = \mathcal{O}_P \left(\sqrt{b} f_n(\Delta) \right) + o_P(1),$$

for $w_1 = (1, 0, 0)$, $w_2 = (0, 1, 0)$, and $w_3 = (0, 0, 1)$. Moreover, since D has non-empty interior, there exists constants $c_1, c_2 > 0$ such that $c_1 f_n(\Delta) \leq f_n(D) \leq c_2 f_n(\Delta)$. Thus, the condition in the Theorem implies $\sqrt{n} f_n(\Delta) \rightarrow \infty$, and the lower bound (11) yields $T_n^*(D) \geq \tilde{c} \sqrt{n} f_n(\Delta) \rightarrow \infty$. On the other hand, $a_n^*(\alpha, D) \leq \tilde{C} \sqrt{b} f_n(\Delta) + 1 \ll T_n^*(D)$. This establishes (7). \square

Funding

The authors acknowledge support from Deutsche Forschungsgemeinschaft (DFG, grant STE 1034/11-2).

References

- ANDREWS, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation. *Econometrica* **59** 817. <https://doi.org/10.2307/2938229>. MR1106513
- BICKEL, P. J. and LEVINA, E. (2008a). Covariance Regularization by Thresholding. *Annals of Statistics* **36** 2577–2604. MR2485008
- BICKEL, P. J. and LEVINA, E. (2008b). Regularized Estimation of Large Covariance Matrices. *The Annals of Statistics* **36**. <https://doi.org/10.1214/009053607000000758>. MR2387969
- BODNAR, T., GUPTA, A. K. and PAROLYA, N. (2014). On the strong convergence of the optimal linear shrinkage estimator for large dimensional covariance matrix. *Journal of Multivariate Analysis* **132** 215–228. <https://doi.org/10.1016/j.jmva.2014.08.006>. MR3266272
- BOURS, M. and STELAND, A. (2021). Large-Sample Approximations and Change Testing for High-Dimensional Covariance Matrices of Multivariate Linear Time Series and Factor Models. *Scandinavian Journal of Statistics* **48** 610–654. <https://doi.org/10.1111/sjos.12508>. MR4275159
- CAI, T. T., REN, Z. and ZHOU, H. H. (2013). Optimal Rates of Convergence for Estimating Toeplitz Covariance Matrices. *Probability Theory and Related Fields* **156** 101–143. <https://doi.org/10.1007/s00440-012-0422-7>. MR3055254
- CAI, T. T., REN, Z. and ZHOU, H. H. (2016). Estimating Structured High-Dimensional Covariance and Precision Matrices: Optimal Rates and Adaptive Estimation. *Electronic Journal of Statistics* **10** 1–59. <https://doi.org/10.1214/15-EJS1081>. MR3466172
- CAI, T. T., ZHANG, C. H. and ZHOU, H. H. (2010). Optimal Rates of Convergence for Covariance Matrix Estimation. *Annals of Statistics* **38** 2118–2144. <https://doi.org/10.1214/09-AOS752>. MR2676885
- DETTE, H., WU, W. and ZHOU, Z. (2019). Change Point Analysis of Correlation in Non-stationary Time Series. *Statistica Sinica* **29** 611–643. <https://doi.org/10.5705/ss.202016.0493>. MR3931381
- DONOHO, D., GAVISH, M. and JOHNSTONE, I. (2018). Optimal shrinkage of eigenvalues in the spiked covariance model. *The Annals of Statistics* **46** 1742–1778. <https://doi.org/10.1214/17-AOS1601>. MR3819116
- ELDAN, R., MIKULINCER, D. and ZHAI, A. (2020). The CLT in High Dimensions: Quantitative Bounds via Martingale Embedding. *The Annals of Probability* **48** 2494–2524. <https://doi.org/10.1214/20-AOP1429>. MR4152649
- GÓRECKI, T., HORVÁTH, L. and KOKOSZKA, P. (2018). Change Point Detection in Heteroscedastic Time Series. *Econometrics and Statistics* **7** 63–88. <https://doi.org/10.1016/j.ecosta.2017.07.005>. MR3824127
- JOHNSTONE, I. M. (2001). On the Distribution of the Largest Eigenvalue in Principal Components Analysis. *The Annals of Statistics* **29**. <https://doi.org/10.1214/aos/1009210544>. MR1863961
- JOHNSTONE, I. M. and LU, A. Y. (2009). On Consistency and Sparsity for Principal Components Analysis in High Dimensions. *Journal of the Ameri-*

- can Statistical Association* **104** 682–693. <https://doi.org/10.1198/jasa.2009.0121>. MR2751448
- KARMAKAR, S. and WU, W. B. (2020). Optimal Gaussian Approximation For Multiple Time Series. *Statistica Sinica* **30** 1399–1417. <https://doi.org/10.5705/ss.202017.0303>. MR4257539
- KOMLÓS, J., MAJOR, P. and TUSNÁDY, G. (1975). An Approximation of Partial Sums of Independent RV'-s, and the Sample DF. I. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete* **32** 111–131. <https://doi.org/10.1007/BF00533093>. MR0375412
- LEDOIT, O. and WOLF, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance* **10** 603–621. [https://doi.org/10.1016/S0927-5398\(03\)00007-0](https://doi.org/10.1016/S0927-5398(03)00007-0).
- LEDOIT, O. and WOLF, M. (2004). A Well-Conditioned Estimator for Large-Dimensional Covariance Matrices. *Journal of Multivariate Analysis* **88** 365–411. [https://doi.org/10.1016/S0047-259X\(03\)00096-4](https://doi.org/10.1016/S0047-259X(03)00096-4). MR2026339
- MAUSE, N. and STELAND, A. (2020). Detecting Changes in the Second Moment Structure of High-Dimensional Sensor-Type Data in a K-Sample Setting. *Sequential Analysis* **39** 336–366. <https://doi.org/10.1080/07474946.2020.1823192>. MR4196430
- MIES, F. (2021). Functional Estimation and Change Detection for Nonstationary Time Series. *Journal of the American Statistical Association* **to appear**. <https://doi.org/10.1080/01621459.2021.1969239>. MR4595473
- MIES, F. and STELAND, A. (2023). Sequential Gaussian Approximation for Nonstationary Time Series in High Dimensions. *Bernoulli* **29** 3114–3140. MR4632133
- NEWKEY, W. K. and WEST, K. D. (1994). Automatic Lag Selection in Covariance Matrix Estimation. *The Review of Economic Studies* **61** 631–653. <https://doi.org/10.2307/2297912>. MR1299308
- PEŠTA, M. and WENDLER, M. (2020). Nuisance-Parameter-Free Changepoint Detection in Non-Stationary Series. *TEST* **29** 379–408. <https://doi.org/10.1007/s11749-019-00659-1>. MR4095034
- SANCETTA, A. (2008). Sample covariance shrinkage for high dimensional dependent data. *J. Multivariate Anal.* **99** 949–967. <https://doi.org/10.1016/j.jmva.2007.06.004>. MR2405100
- SCHMIDT, S. K., WORNOWIZKI, M., FRIED, R. and DEHLING, H. (2021). An asymptotic test for constancy of the variance under short-range dependence. *The Annals of Statistics* **49** 3460–3481. MR4352537
- STELAND, A. (2018). Shrinkage for Covariance Estimation: Asymptotics, Confidence Intervals, Bounds and Applications in Sensor Monitoring and Finance. *Statistical Papers* **59** 1441–1462. <https://doi.org/10.1007/s00362-018-1040-y>. MR3878741
- STELAND, A. (2020). Testing and Estimating Change-Points in the Covariance Matrix of a High-Dimensional Time Series. *Journal of Multivariate Analysis* **177** 104582. <https://doi.org/10.1016/j.jmva.2019.104582>.

- MR4068598
- STELAND, A. (2023). Online Detection of Changes in Moment-Based Projections: When to Retrain Deep Learners or Update Portfolios? *arXiv:2302.07198*.
- STELAND, A. (2024). Flexible Nonlinear Inference and Change-Point Testing of High-Dimensional Spectral Density Matrices. *Journal of Multivariate Analysis* **199** 105245. <https://doi.org/10.1016/j.jmva.2023.105245>. MR4666257
- STELAND, A. and VON SACHS, R. (2017). Large-Sample Approximations for Variance-Covariance Matrices of High-Dimensional Time Series. *Bernoulli* **23** 2299–2329. <https://doi.org/10.3150/16-BEJ811>. MR3648032
- STELAND, A. and VON SACHS, R. (2018). Asymptotics for High-Dimensional Covariance Matrices and Quadratic Forms with Applications to the Trace Functional and Shrinkage. *Stochastic Processes and their Applications* **128** 2816–2855. <https://doi.org/10.1016/j.spa.2017.10.007>. MR3811705
- SWADE, A., NOLTE, S., SHACKLETON, M. and LOHRE, H. (2023). Why do equally weighted portfolios beat value-weighted ones? *Journal of Portfolio Management* **49** 167–187.
- WU, W. B. (2005). Nonlinear System Theory: Another Look at Dependence. *Proceedings of the National Academy of Sciences* **102** 14150–14154. <https://doi.org/10.1073/pnas.0506715102>. MR2172215
- WU, W. B. and ZHOU, Z. (2011). Gaussian Approximations for Non-Stationary Multiple Time Series. *Statistica Sinica* **21** 1397–1413. <https://doi.org/10.5705/ss.2008.223>. MR2827528
- ZEILEIS, A., KÖLL, S. and GRAHAM, N. (2020). Various Versatile Variances: An Object-Oriented Implementation of Clustered Covariances in R. *Journal of Statistical Software* **95** 1–36. <https://doi.org/10.18637/jss.v095.i01>.
- ZHOU, Z. (2013). Heteroscedasticity and Autocorrelation Robust Structural Change Detection. *Journal of the American Statistical Association* **108** 726–740. <https://doi.org/10.1080/01621459.2013.787184>. MR3174655