

Adaptive estimation of the copula correlation matrix for semiparametric elliptical copulas

MARTEN WEGKAMP¹ and YUE ZHAO²

¹*Department of Mathematics and Department of Statistical Science, Cornell University, Ithaca, NY 14853, USA. E-mail: marten.wegkamp@cornell.edu*

²*Department of Statistical Science, Cornell University, Ithaca, NY 14853, USA. E-mail: yz453@cornell.edu*

We study the adaptive estimation of copula correlation matrix Σ for the semi-parametric elliptical copula model. In this context, the correlations are connected to Kendall's tau through a sine function transformation. Hence, a natural estimate for Σ is the plug-in estimator $\widehat{\Sigma}$ with Kendall's tau statistic. We first obtain a sharp bound on the operator norm of $\widehat{\Sigma} - \Sigma$. Then we study a factor model of Σ , for which we propose a refined estimator $\widetilde{\Sigma}$ by fitting a low-rank matrix plus a diagonal matrix to $\widehat{\Sigma}$ using least squares with a nuclear norm penalty on the low-rank matrix. The bound on the operator norm of $\widehat{\Sigma} - \Sigma$ serves to scale the penalty term, and we obtain finite sample oracle inequalities for $\widetilde{\Sigma}$. We also consider an elementary factor copula model of Σ , for which we propose closed-form estimators. All of our estimation procedures are entirely data-driven.

Keywords: correlation matrix; elliptical copula; factor model; Kendall's tau; nuclear norm regularization; oracle inequality; primal-dual certificate

1. Introduction

1.1. Background

A popular model for high dimensional data is the *semi-parametric elliptical copula model* [13, 23,24,29], the family of distributions whose dependence structures are specified by parametric elliptical copulas but whose marginal distributions are left unspecified. The elliptical copula of a d -variate distribution from the semi-parametric elliptical copula model is uniquely characterized by a *characteristic generator* ϕ and a *copula correlation matrix* $\Sigma \in \mathbb{R}^{d \times d}$. We refer the readers to Appendix A for a more detailed discussion about these concepts. For simplicity of presentation, we will make the blanket assumption that all random vectors we consider have continuous marginals.

The semi-parametric elliptical copula model includes numerous families of distributions of popular interest. For instance, we recover from this model distributions with Gaussian copulas, sometimes referred to in recent literature as the nonparanormal model [30], by choosing the particular characteristic generator $\phi(t) = \exp(-t/2)$.

Throughout the paper, we assume that the random vector $X \in \mathbb{R}^d$ follows a distribution from the semi-parametric elliptical copula model, and in particular we let X have copula correlation matrix Σ . We let $X^1, \dots, X^n \in \mathbb{R}^d$, with $X^i = (X^i_1, \dots, X^i_d)^T$, be a sequence of independent

copies of X . We recall the formulas for (the population version of) Kendall’s tau between the k th and ℓ th coordinates,

$$\tau_{k\ell} = \mathbb{E}[\text{sgn}(X_k^1 - X_k^2) \text{sgn}(X_\ell^1 - X_\ell^2)], \tag{1.1}$$

and the corresponding Kendall’s tau statistic,

$$\widehat{\tau}_{k\ell} = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} [\text{sgn}(X_k^i - X_k^j) \text{sgn}(X_\ell^i - X_\ell^j)]. \tag{1.2}$$

We let (the population version of) the Kendall’s tau matrix T have entries

$$[T]_{k\ell} = \tau_{k\ell} \quad \text{for all } 1 \leq k, \ell \leq d,$$

and estimate T using the empirical Kendall’s tau matrix \widehat{T} with entries

$$[\widehat{T}]_{k\ell} = \widehat{\tau}_{k\ell} \quad \text{for all } 1 \leq k, \ell \leq d. \tag{1.3}$$

We note that \widehat{T} is a matrix U -statistic because it can be written as

$$\widehat{T} = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} [\text{sgn}(X^i - X^j) \text{sgn}(X^i - X^j)^T].$$

In addition, we note the basic facts that T is the correlation matrix of the centered random vector $\text{sgn}(X^1 - X^2)$ and so in particular is positive semidefinite, that \widehat{T} , as a scaled sum of rank-one positive semidefinite matrices $\text{sgn}(X^i - X^j) \text{sgn}(X^i - X^j)^T$ for $1 \leq i < j \leq n$, is also positive semidefinite, and that $\mathbb{E}[\widehat{T}] = T$.

For the semi-parametric elliptical copula model, we can relate the elements of the copula correlation matrix Σ to the elements of the Kendall’s tau matrix T independently of the characteristic generator via the formula

$$\Sigma = \sin\left(\frac{\pi}{2} T\right); \tag{1.4}$$

see [14,21,22,26,27]. Here and throughout the paper, we use the convention that the sign, sine and cosine functions act component-wise when supplied with a vector or a matrix as their argument; hence equation (1.4) specifies that

$$[\Sigma]_{k\ell} = \sin\left(\frac{\pi}{2} \tau_{k\ell}\right) \quad \text{for all } 1 \leq k, \ell \leq d.$$

This simple and elegant relationship has contributed to the popularity of elliptical distributions and the semi-parametric elliptical copula model, and has led to the widespread application of the *plug-in estimator* $\widehat{\Sigma}$ of Σ given by

$$\widehat{\Sigma} = \sin\left(\frac{\pi}{2} \widehat{T}\right); \tag{1.5}$$

see, for instance, [11,13,23,24,28,47]. Here, we briefly review some recent advances involving the plug-in estimator. [23] studies the property of $\widehat{\Sigma}$ as an estimator of Σ in the *asymptotic* setting with the dimension d fixed under the assumption of an *elliptical copula correlation factor model*, whose precise definition will be introduced later in Section 1.2. For distributions with Gaussian copulas, [28] employs $\widehat{\Sigma}$ to study the estimation of precision matrix, that is, Σ^{-1} , under a sparsity assumption on Σ^{-1} , and a sharp bound on the element-wise ℓ_∞ norm of $\widehat{\Sigma} - \Sigma$ is central to their analysis.¹

1.2. Proposed research

We aim to present in this paper precise estimators of the copula correlation matrix Σ .

In Section 2, we focus on the plug-in estimator $\widehat{\Sigma}$, and present a sharp (upper) bound on the operator norm of $\widehat{\Sigma} - \Sigma$, which we denote by $\|\widehat{\Sigma} - \Sigma\|_2$. To the best of our knowledge, our bound on $\|\widehat{\Sigma} - \Sigma\|_2$ is new, even for distributions with Gaussian copulas. Here, we list some of the potential applications of this bound. First, it has often been observed that the plug-in estimator $\widehat{\Sigma}$ is not always positive semidefinite [11,23]. This not only is a discomfoting problem by itself but also limits the potential application of the plug-in estimator; for example, certain Graphical Lasso algorithms [16] may fail on input that is not positive semidefinite. We refer the readers to [45] for a more detailed discussion and another example involving the Markowitz portfolio optimization problem. Our bound on $\|\widehat{\Sigma} - \Sigma\|_2$ will precisely quantify the extent to which the nonpositive semidefinite problem may happen; for instance, if the smallest eigenvalue of Σ exceeds the bound on $\|\widehat{\Sigma} - \Sigma\|_2$, then $\widehat{\Sigma}$ will be positive definite.

As we were completing this manuscript, we became aware of a result by Fang Han and Han Liu in [17] that is similar to (our) inequality (2.7a) in Theorem 2.2. In deriving their result, they also employed matrix concentration inequalities to arrive at a version of inequality (2.1a); then they invoked different proof techniques to arrive at a version of Lemma 4.3, which led to their version of inequality (2.6). Our work is independent.

A second application of the bound on $\|\widehat{\Sigma} - \Sigma\|_2$ appears in Section 3. Here, we study the elliptical copula correlation factor model, which postulates that the copula correlation matrix Σ of X admits the decomposition

$$\Sigma = \Theta^* + V^* \tag{1.6}$$

for some low-rank or nearly low-rank, positive semidefinite matrix $\Theta^* \in \mathbb{R}^{d \times d}$ and some diagonal matrix $V^* \in \mathbb{R}^{d \times d}$ with nonnegative diagonal entries. In this case, if Θ^* admits the decomposition $\Theta^* = LL^T$ for some $L \in \mathbb{R}^{d \times r}$, then there exists elliptically distributed $\xi \in \mathcal{E}_{r+d}(0, I_{r+d}, \phi)$ (here we invoke the notation of Definition A.1) for the $(r+d) \times (r+d)$ identity

¹We note that, under the setting of distributions with Gaussian copulas, analogous to equation (1.4), we also have $\Sigma = 2\sin((\pi/6)R)$ for R the matrix of (the population version of) Spearman's rho. Inspired by this observation, both [28] and [46] employ $\widehat{\Sigma}^\rho$, a variant of $\widehat{\Sigma}$ using Spearman's rho statistic, to study the estimation of precision matrix under this setting. In contrast to Kendall's tau, however, once we generalize from distributions with Gaussian copulas to the semi-parametric elliptical copula model, Spearman's rho is no longer invariant within the family of distributions with the same copula correlation matrix [21], that is, a simple relationship analogous to equation (1.4) ceases to exist for Spearman's rho in this wider context. Hence, we do not pursue an estimation procedure using Spearman's rho.

matrix I_{r+d} and some characteristic generator ϕ such that X and $(L, V^{*1/2})\xi$ have the same copula. Here, we note that the components of ξ are merely uncorrelated, instead of independent as in the case for standard factor analysis where normality is assumed. Consideration of the potential dimension reduction offered by the factor model and the fact that the diagonal elements of the target copula correlation matrix Σ are all equal to one leads us to propose a refined estimator $\tilde{\Sigma}$ of Σ . In short, we fit the off-diagonal elements of a low-rank matrix to the off-diagonal elements of $\widehat{\Sigma}$ using least squares with a nuclear norm penalty on the low-rank matrix; then we obtain the refined estimator $\tilde{\Sigma}$ from the low-rank matrix by setting the diagonal elements of the latter to one. The bound on $\|\widehat{\Sigma} - \Sigma\|_2$ will serve to scale the penalty term. As we will discuss in detail in Section 3.3, our problem is a variant of the matrix completion problem, but in contrast to the existing literature, the special diagonal structure of V^* enables us to perform much more precise analysis. In the end, our oracle inequality for $\tilde{\Sigma}$ holds under a single, very mild condition on the low-rank component Θ^* , and balances the approximation error with the estimation error, with the latter roughly proportional to the number of parameters in the model divided by the sample size.

As a warm-up to the general setting above, we will also consider the *elementary factor copula model*, a special instance of the elliptical copula correlation factor model in which V^* is proportional to the $d \times d$ identity matrix I_d . For this model, we will propose and study closed-form estimators.

Throughout our studies, we will provide entirely data-driven estimation procedures involving explicit constants and measurable quantities. In addition, we will establish positive semidefinite versions of the plug-in estimator, the closed-form estimator and the refined estimator of the copula correlation matrix, with minimal loss in performance.

1.3. Notation

For any matrix A , we will use $[A]_{k\ell}$ to denote the k, ℓ th element of A (i.e., the entry on the k th row and ℓ th column of A). For a vector $x \in \mathbb{R}^m$, we denote by $\text{diag}^*(x) \in \mathbb{R}^{m \times m}$ the diagonal matrix with $[\text{diag}^*(x)]_{ii} = x_i$ for $i = 1, \dots, m$. We let the constant α with $0 < \alpha < 1$ be arbitrary, but typically small; we will normally bound stochastic events with probability at least $1 - \mathcal{O}(\alpha)$. We let I_d denote the identity matrix in $\mathbb{R}^{d \times d}$. In this paper, the majority of the vectors will belong to \mathbb{R}^d , and the majority of the matrices will be symmetric and belong to $\mathbb{R}^{d \times d}$; notable exceptions to the latter rule include some matrices of left or right singular vectors. For notational brevity, we will not always explicitly specify the dimension of a matrix when such information could be inferred from the context. The Frobenius inner product $\langle \cdot, \cdot \rangle$ on the space of matrices is defined as $\langle A, B \rangle = \text{tr}(A^T B)$ for commensurate matrices A, B . For norms on matrices, we use $\|\cdot\|_2$ to denote the operator norm, $\|\cdot\|_*$ the nuclear norm (i.e., the sum of singular values), $\|\cdot\|_F$ the Frobenius norm resulting from the Frobenius inner product, $\|\cdot\|_\infty$ the element-wise ℓ_∞ norm (i.e., $\|A\|_\infty = \max_{k,\ell} |[A]_{k\ell}|$), and $\|\cdot\|_1$ the element-wise ℓ_1 norm. The effective rank of a positive semidefinite matrix A is defined as $r_e(A) = \text{tr}(A)/\|A\|_2$. We let $\lambda_{\max}(\cdot)$ and $\lambda_{\min}(\cdot)$ denote the largest and the smallest eigenvalues, respectively, and let \mathcal{S}_+^d be the set of $d \times d$ correlation matrices, that is, positive semidefinite matrices with all diagonal elements equal to one. We use \circ to denote the Hadamard (or Schur) product. For notational brevity when studying

the factor model, for an arbitrary matrix $A \in \mathbb{R}^{d \times d}$, we let $A_o \in \mathbb{R}^{d \times d}$ be the matrix with the same off-diagonal elements as A , but with all diagonal elements equal to zero, that is,

$$A_o = A - I_d \circ A. \tag{1.7}$$

Again for notational brevity, this time when establishing probability bounds involving Kendall's tau statistics, we will assume throughout that the number of samples, n , is even, and denote

$$f(n, d, \alpha) = \sqrt{\frac{16}{3} \cdot \frac{d \cdot \log(2\alpha^{-1}d)}{n}}.$$

Remark. When n is odd, the appropriate f to use is

$$f(n, d, \alpha) = \sqrt{\frac{16}{3} \cdot \frac{d \cdot \log(2\alpha^{-1}d)}{2\lfloor n/2 \rfloor}}.$$

This is due to the fact that when n is odd, we can group X^1, \dots, X^n into at most $\lfloor n/2 \rfloor$ pairs of (X^i, X^j) 's such that the different pairs are independent.

2. Plug-in estimation of the copula correlation matrix

In this section, we focus on the plug-in estimator $\widehat{\Sigma}$ of the copula correlation matrix Σ and in particular provide a bound on $\|\widehat{\Sigma} - \Sigma\|_2$. We recall that Σ is related to the Kendall's tau matrix T via a sine function transformation as in equation (1.4), and $\widehat{\Sigma}$ is related to the empirical Kendall's tau matrix \widehat{T} via the same transformation as in equation (1.5). We note that a typical proof for a bound on $\|\widehat{\Sigma} - \Sigma\|_\infty$ in the existing literature first establishes a bound on $\|\widehat{T} - T\|_\infty$ through a combination of Hoeffding's classical bound for the (scalar) U -statistic applied to each element of $\widehat{T} - T$ and a union bound argument, and then establishes the bound on $\|\widehat{\Sigma} - \Sigma\|_\infty$ through the Lipschitz property of the sine function transformation [28]. Our proof for the bound on $\|\widehat{\Sigma} - \Sigma\|_2$ is similarly divided into two essentially independent stages:

1. First, in Section 2.1, we establish a bound on $\|\widehat{T} - T\|_2$. This stage can be considered as the matrix counterpart in terms of the operator norm to Hoeffding's classical bound for the (scalar) U -statistic;
2. Next, in Section 2.2, we bound $\|\widehat{\Sigma} - \Sigma\|_2$ by a constant times $\|\widehat{T} - T\|_2$ up to an additive quadratic term in $f(n, d, \alpha)$. This stage can be considered as the matrix counterpart in terms of the operator norm to the Lipschitz property of the sine function transformation. Then, combined with the bound on $\|\widehat{T} - T\|_2$, we establish the bound on $\|\widehat{\Sigma} - \Sigma\|_2$.

2.1. Bounding $\|\widehat{T} - T\|_2$

In this section, we bound $\|\widehat{T} - T\|_2$, establishing both data-driven and data-independent versions. We rely on the results from [42] out of the vast literature on matrix concentration inequalities (see [4,43] for a glimpse of the literature).

Theorem 2.1. *We have, with probability at least $1 - \alpha$,*

$$\|\widehat{T} - T\|_2 < \max\{\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)\} \tag{2.1a}$$

$$\leq \sqrt{\|\widehat{T}\|_2 f^2(n, d, \alpha) + \frac{1}{4} f^4(n, d, \alpha) + \frac{1}{2} f^2(n, d, \alpha)} \tag{2.1b}$$

$$< \max\{\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)\} + f^2(n, d, \alpha). \tag{2.1c}$$

Remark. By decoupling the matrix U -statistic $\widehat{T} - T$ using (4.5), and [42], inequality (6.1.3) in Theorem 6.1.1, we can also obtain a bound on $\mathbb{E}[\|\widehat{T} - T\|_2]$. We omit the details.

Proof of Theorem 2.1. The proof can be found in Section 4. □

We elaborate the results presented in Theorem 2.1. First, we note that the bound offered by inequality (2.1a) is the tightest, but contains the possibly unknown population quantity $\|T\|_2$. Hence, we also derive a data-driven bound (2.1b), whose performance is in turn guaranteed by (2.1c) in terms of the deterministic $\|T\|_2$. Theorem 2.1 also shows that the right-hand side of (2.1b) is no more than $f^2(n, d, \alpha)$ away from the right-hand side of (2.1a). This is because the former is sandwiched between the right-hand sides of (2.1a) and (2.1c), and the latter two terms differ by $f^2(n, d, \alpha)$.

Next, for latter convenience, we note that when n is large enough such that

$$\|T\|_2 \geq f^2(n, d, \alpha) = \frac{16}{3} \cdot \frac{d \cdot \log(2\alpha^{-1}d)}{n}, \tag{2.2}$$

the first term dominates the second term in the curly bracket on the right-hand side of (2.1a), that is,

$$\max\{\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)\} = \sqrt{\|T\|_2} f(n, d, \alpha). \tag{2.3}$$

Finally, we discuss the optimality of Theorem 2.1, specifically inequality (2.1a). First, we compare our result to some recent upper bounds established by other authors under conditions related to but more restrictive than the semi-parametric elliptical copula model. Under the same model but with the additional “sign sub-Gaussian condition,” [17] establishes in their Theorem 4.10 that

$$\|\widehat{T} - T\|_2 = \mathcal{O}\left(\|T\|_2 \sqrt{\frac{d + \log(\alpha^{-1})}{n}}\right) \tag{2.4}$$

with probability at least $1 - 2\alpha$. Meanwhile, for distributions with Gaussian copulas, [33] establishes in their Corollary 3 a more complicated bound which, in the regime $n \geq d$, $\|T\|_2 \geq \max\{\log(d), \log(\alpha^{-1})\}$ and $\|\Sigma\|_{2, \max} \leq \|\Sigma\|_2^{1/2}$, reduces to that inequality (2.4) holds with probability at least $1 - \alpha$. Here $\|\Sigma\|_{2, \max} = \max_{\|u\|=1} \|\Sigma u\|_{\max}$ with $\|\cdot\|$ and $\|\cdot\|_{\max}$ being the Euclidean norm and the element-wise ℓ_∞ norm for vectors, respectively.

Such bounds, which are based on Gaussian concentration inequalities, are of a different flavor. Nevertheless, here we will attempt a very crude comparison. We set $\alpha = 1/d$ so that both our

inequality (2.1a) and inequality (2.4) hold with probability at least $1 - \mathcal{O}(1/d)$. We also assume that n is large enough such that inequality (2.2) holds. Then the right-hand sides of (2.1a) and (2.4) are $\mathcal{O}(\sqrt{\|T\|_2 d \log(d)/n})$ and $\mathcal{O}(\|T\|_2 \sqrt{d/n})$, respectively. Hence, the bound provided by our inequality (2.1a) sheds an operator norm factor $\sqrt{\|T\|_2}$ at the expense of an extra log factor $\sqrt{\log(d)}$.

From another angle, we contrast our upper bound (2.1a) to the corresponding lower bound implied by the argument presented in the proof of [31], Theorem 2, in the context of covariance matrix estimation. Such a comparison reveals that our bound (2.1a) is optimal up to the (aforementioned) operator norm factor $\sqrt{\|T\|_2}$ and the log factor $\sqrt{\log(d)}$ in $f(n, d, \alpha)$. The study of if and when these factors can be removed is beyond the scope of this paper.² We also note that, by [42], Chapter 7, in inequality (2.1a), we could replace the ambient dimension d inside the log function in $f(n, d, \alpha)$ by $\tilde{d} = 4d/\|T\|_2$. Here, \tilde{d} is the effective rank of a semidefinite upper bound of $\mathbb{E}[(\hat{T} - T)^2]$ with \hat{T} defined in equation (4.1). Hence, if $\|T\|_2$ is comparable to d , then the log factor is effectively removed. In large sample size or large dimension setting, it is customary to set α to be $1/\max\{n, d\}$ so that the exclusion probability α tends to zero as n or d increases. For such a setting of α , we shed at most a constant multiplicative factor in the bound on $\|\hat{T} - T\|_2$ by setting d to \tilde{d} inside the log function. Thus, for brevity of presentation in later sections, we have avoided invoking the effective rank.

2.2. Bounding $\|\hat{\Sigma} - \Sigma\|_2$ in terms of $\|\hat{T} - T\|_2$

In this section, we establish in Theorem 2.2 the promised link between $\|\hat{\Sigma} - \Sigma\|_2$ and $\|\hat{T} - T\|_2$. Based on this result, we establish bounds on $\|\hat{\Sigma} - \Sigma\|_2$ in the same theorem.

We also establish in Theorem 2.2 a link between $\|\hat{T}' - T\|_2$ and $\|\hat{\Sigma}' - \Sigma\|_2$, for \hat{T}' that is any generic estimator of T (i.e., \hat{T}' is not necessarily the empirical Kendall's tau matrix \hat{T}), and $\hat{\Sigma}'$ the resulting generic plug-in estimator, that is,

$$\hat{\Sigma}' = \sin\left(\frac{\pi}{2} \hat{T}'\right).$$

Possibilities of generic estimators \hat{T}' of T include regularized estimators such as thresholding [2,6] or tapering [5] estimator. Such generic estimators \hat{T}' of T and the resulting generic plug-in estimators $\hat{\Sigma}'$ of Σ have the potential to provide faster convergence rate than the empirical Kendall's tau matrix \hat{T} and the plug-in estimator $\hat{\Sigma}$ if appropriate structure of T is known in advance so a regularized estimator \hat{T}' could be used. Hence, we briefly include the consideration of generic estimators in Theorem 2.2.

An auxiliary result relating $\|T\|_2$ to $\|\Sigma\|_2$ is provided by Theorem 2.3.

²By our proof of Theorem 2.1, inequality (2.1a) also holds with the replacement of \hat{T} by its decoupled version \tilde{T} defined in (4.1). Then, by the argument of [42], Section 6.1.2, we can show that the operator norm factor $\sqrt{\|T\|_2}$ is in fact necessary in this variant of (2.1a) in terms of \tilde{T} at least in certain scenarios. Unfortunately, the same argument does not apply directly to (2.1a) in terms of the matrix U -statistic \hat{T} .

Theorem 2.2. Let \widehat{T}' be a generic estimator of T , and $\widehat{\Sigma}'$ the resulting generic plug-in estimator of Σ . We have, for some absolute constants C'_1, C'_2 (we may take $C'_1 = \pi$ and $C'_2 = \pi^2/8 < 1.24$),

$$\|\widehat{\Sigma}' - \Sigma\|_2 \leq C'_1 \|\widehat{T}' - T\|_2 + C'_2 \|\widehat{T}' - T\|_2^2. \quad (2.5)$$

Recall \widehat{T} as defined in equation (1.3) and the resulting plug-in estimator $\widehat{\Sigma}$ as defined in equation (1.5). We have, for some absolute constants C_1, C_2 (we may take $C_1 = \pi$ and $C_2 = 3\pi^2/16 < 1.86$), with probability at least $1 - \frac{1}{4}\alpha^2$,

$$\|\widehat{\Sigma} - \Sigma\|_2 \leq C_1 \|\widehat{T} - T\|_2 + C_2 f^2(n, d, \alpha). \quad (2.6)$$

Recall that Theorem 2.1 bounds $\|\widehat{T} - T\|_2$. Hence, starting from inequality (2.6), we have, with probability at least $1 - \alpha - \frac{1}{4}\alpha^2$,

$$\|\widehat{\Sigma} - \Sigma\|_2 < C_1 \max\{\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)\} + C_2 f^2(n, d, \alpha) \quad (2.7a)$$

$$\leq C_1 \sqrt{\|\widehat{T}\|_2 f^2(n, d, \alpha) + \frac{1}{4} f^4(n, d, \alpha)} + (\frac{1}{2}C_1 + C_2) f^2(n, d, \alpha) \quad (2.7b)$$

$$< C_1 \max\{\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)\} + (C_1 + C_2) f^2(n, d, \alpha). \quad (2.7c)$$

Proof. The proof can be found in Section 4. □

We elaborate the results presented in Theorem 2.2. First, the relationship between the bounds (2.7a), (2.7b) and (2.7c) is analogous to the relationship between the bounds (2.1a), (2.1b) and (2.1c) as has been discussed following Theorem 2.1. Next, we discuss the relative merits of inequalities (2.5) and (2.6). We note that

1. For the plug-in estimator $\widehat{\Sigma}$, instead of starting from inequality (2.6), we can also start from inequality (2.5), take the particular choices $\widehat{T}' = \widehat{T}$ and $\widehat{\Sigma}' = \widehat{\Sigma}$, and establish a bound on $\|\widehat{\Sigma} - \Sigma\|_2$ via inequality (2.1a) in Theorem 2.1 as

$$\|\widehat{\Sigma} - \Sigma\|_2 \leq \max\{C'_1 \sqrt{\|T\|_2} f(n, d, \alpha) + C'_2 \|T\|_2 f^2(n, d, \alpha), C'_1 f^2(n, d, \alpha) + C'_2 f^4(n, d, \alpha)\}$$

with probability at least $1 - \alpha$. However, it is obvious that this bound is not as tight as the one presented in inequality (2.7a), which we obtained via inequality (2.6).

2. On the other hand, suppose that we have a generic plug-in estimator $\widehat{\Sigma}'$ of Σ based on a generic estimator \widehat{T}' of T that achieves a rate $\|\widehat{T}' - T\|_2 \ll f(n, d, \alpha)$ (a rate faster than the one for $\|\widehat{T} - T\|_2$). Then, inequality (2.5) would yield

$$\|\widehat{\Sigma}' - \Sigma\|_2 \ll C'_1 f(n, d, \alpha) + C'_2 f^2(n, d, \alpha),$$

which is tighter than the bound offered by inequality (2.7a).

Therefore, whether inequality (2.5) or (2.6) should be preferred depends on the available estimator of T and the rate of convergence of the estimator.

Inequalities (2.7a) and (2.7c) in Theorem 2.2 contain the term $\|T\|_2$. Using the result of Theorem 2.3, we could relate $\|T\|_2$ back to $\|\Sigma\|_2$, so that we bound $\|\widehat{\Sigma} - \Sigma\|_2$ directly in terms of the copula correlation matrix Σ .

Theorem 2.3. *We have*

$$\frac{2}{\pi} \|\Sigma\|_2 \leq \|T\|_2 \leq \|\Sigma\|_2. \tag{2.8}$$

Hence, inequalities (2.7a) and (2.7c) hold with $\|T\|_2$ replaced by $\|\Sigma\|_2$.

Remark. The second half of inequality (2.8) is tight: $\|T\|_2 = \|\Sigma\|_2$ when $T = \Sigma = I_d$.

Proof of Theorem 2.3. The proof can be found in Section 4. □

2.3. Obtaining a positive semidefinite estimator $\widehat{\Sigma}^+$ from the plug-in estimator $\widehat{\Sigma}$

As has been mentioned in Section 1.2, the plug-in estimator $\widehat{\Sigma}$ may fail to be positive semidefinite. In this section, we demonstrate a procedure that, in such an event, obtains an explicitly positive semidefinite estimator $\widehat{\Sigma}^+$ of Σ from $\widehat{\Sigma}$ with minimal loss in performance. The procedure is suggested by a referee and is inspired by [45]. Note that, when $\widehat{\Sigma}$ is not positive semidefinite, we cannot simply set all the negative eigenvalues of $\widehat{\Sigma}$ to zero, because the resulting estimator will still not be a correlation matrix, specifically because some of the diagonal elements of the resulting estimator will exceed one.

In order to also cover the closed-form estimator and the refined estimator when we study a factor model for Σ , we will consider a more general situation. We let $\|\cdot\|$ be a generic matrix norm and $\widehat{\Sigma}^{\text{generic}}$ a generic estimator of Σ . We do not require $\widehat{\Sigma}^{\text{generic}}$ to be a correlation matrix. We let the feasible region $\mathcal{F} \subset \mathbb{R}^{d \times d}$ be such that \mathcal{F} is nonempty, closed and convex, satisfies $\mathcal{F} \subset \mathcal{S}_+^d$, but is otherwise arbitrary at this stage. From $\widehat{\Sigma}^{\text{generic}}$, we construct an estimator $\widehat{\Sigma}^{\text{generic}+}$ as

$$\widehat{\Sigma}^{\text{generic}+} = \arg \min_{\Sigma' \in \mathcal{F}} \|\Sigma' - \widehat{\Sigma}^{\text{generic}}\|. \tag{2.9}$$

We note that a solution to the right-hand side of (2.9) always exists. If the norm $\|\cdot\|$ is strictly convex (which is the case for the Frobenius norm), the solution $\widehat{\Sigma}^{\text{generic}+}$ is uniquely determined, while if multiple solutions to the right-hand side of (2.9) exist, we arbitrarily choose one of the solutions to be $\widehat{\Sigma}^{\text{generic}+}$. By construction, $\widehat{\Sigma}^{\text{generic}+}$ is a correlation matrix and so in particular is positive semidefinite. In addition, Theorem 2.4 shows that, when $\Sigma \in \mathcal{F}$, the performance of $\widehat{\Sigma}^{\text{generic}+}$ is comparable to the performance of $\widehat{\Sigma}^{\text{generic}}$ as measured by the deviation from Σ in the norm $\|\cdot\|$.

Theorem 2.4. *Suppose that $\Sigma \in \mathcal{F}$. Then the estimator $\widehat{\Sigma}^{\text{generic}+}$ in (2.9) satisfies*

$$\|\widehat{\Sigma}^{\text{generic}+} - \Sigma\| \leq 2\|\widehat{\Sigma}^{\text{generic}} - \Sigma\|.$$

Proof. The proof can be found in Section 4. □

Theorem 2.4 enables us to obtain from the plug-in estimator $\widehat{\Sigma}$ a positive semidefinite estimator $\widehat{\Sigma}^+$ of Σ such that $\|\widehat{\Sigma}^+ - \Sigma\|_2$ is comparable to $\|\widehat{\Sigma} - \Sigma\|_2$ and, if necessary, $\|\widehat{\Sigma}^+ - \Sigma\|_\infty$

is comparable to $\|\widehat{\Sigma} - \Sigma\|_\infty$, as we demonstrate in Corollary 2.5. As we have mentioned in Section 1.1, a sharp bound on the element-wise ℓ_∞ norm is central in some existing procedures for estimating the precision matrix Σ^{-1} .

Corollary 2.5. *In (2.9), we let the generic matrix norm $\|\cdot\|$ be replaced by the operator norm $\|\cdot\|_2$, the generic estimator $\widehat{\Sigma}^{\text{generic}}$ be replaced by the plug-in estimator $\widehat{\Sigma}$, and the solution $\widehat{\Sigma}^{\text{generic}+}$ be replaced by $\widehat{\Sigma}^+$. First, we choose $\mathcal{F} = \mathcal{S}_+^d$. Then, $\widehat{\Sigma}^+$ satisfies*

$$\|\widehat{\Sigma}^+ - \Sigma\|_2 \leq 2\|\widehat{\Sigma} - \Sigma\|_2. \quad (2.10)$$

Alternatively, we choose $C_3 = \sqrt{3\pi^2/8} < 1.93$, and

$$\mathcal{F} = \{\Sigma': \Sigma' \in \mathcal{S}_+^d \text{ and } \|\Sigma' - \widehat{\Sigma}\|_\infty \leq C_3 d^{-1/2} f(n, d, \alpha)\}. \quad (2.11)$$

Then, with probability at least $1 - \frac{1}{4}\alpha^2$, $\widehat{\Sigma}^+$ satisfies inequality (2.10) and

$$\|\widehat{\Sigma}^+ - \Sigma\|_\infty \leq 2C_3 d^{-1/2} f(n, d, \alpha) \quad (2.12)$$

simultaneously. We recall that $\|\widehat{\Sigma} - \Sigma\|_2$ is bounded as in Theorem 2.2.

Proof. The proof can be found in Section 4. □

3. Estimating the copula correlation matrix in the factor model

In this section, we assume an elliptical copula correlation factor model for $X \in \mathbb{R}^d$. Recall that, under this assumption, the copula correlation matrix Σ of X can be written as

$$\Sigma = \Theta^* + V^*$$

as in equation (1.6), with $\Theta^* \in \mathbb{R}^{d \times d}$ a low-rank or nearly low-rank positive semidefinite³ matrix, and $V^* \in \mathbb{R}^{d \times d}$ a diagonal matrix with nonnegative diagonal entries. Our goal of this section is to present estimators that take advantage of the potential dimension reduction offered by the factor model and the special diagonal structure of V^* .

As a prelude to the main result of this section, in Section 3.1, we first consider the elementary factor copula model, for which we study closed-form estimators. Sections 3.2 and 3.3 form an integral part: in the former, we introduce additional notation, while in the latter we present our main result of Section 3, specifically by constructing the refined estimator $\widetilde{\Sigma}$ of Σ based on the plug-in estimator $\widehat{\Sigma}$ and establishing its associated oracle inequality.

³The case that Θ^* is not positive semidefinite, though unnatural because in the factor model Θ^* should equal LL^T for some matrix L , can be easily accommodated. We restrict our argument to positive semidefinite matrices only to take advantage of the notational brevity offered by the fact that their singular value decomposition and eigen-decomposition coincide.

3.1. Analysis of closed-form estimators in the elementary factor copula model

The elementary factor copula model assumes that $\Theta^* \in \mathbb{R}^{d \times d}$ is a positive semidefinite matrix of unknown rank r with positive eigenvalues $\lambda_1(\Theta^*) \geq \dots \geq \lambda_r(\Theta^*)$, and

$$V^* = \sigma^2 I_d \tag{3.1}$$

with $\sigma^2 > 0$. In other words, the copula correlation matrix Σ admits the decomposition

$$\Sigma = \Theta^* + \sigma^2 I_d.$$

Comparison of the eigen-decomposition

$$\Theta^* + \sigma^2 I_d = U \text{diag}^*(\lambda_1(\Theta^*) + \sigma^2, \dots, \lambda_r(\Theta^*) + \sigma^2, \sigma^2, \dots, \sigma^2) U^T$$

of Σ , with the eigen-decomposition $\sum_{k=1}^d \hat{\lambda}_k \hat{u}_k \hat{u}_k^T$ (with $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_d$) of the plug-in estimator $\hat{\Sigma}$, leads us to propose the following closed-form estimators:

$$\begin{aligned} \hat{r} &= \sum_{k=1}^d \mathbb{1}\{\hat{\lambda}_k - \hat{\lambda}_d \geq \mu\}, \\ \hat{\sigma}^2 &= \frac{1}{d - \hat{r}} \sum_{k > \hat{r}} \hat{\lambda}_k, \\ \hat{\Theta} &= \sum_{k=1}^{\hat{r}} (\hat{\lambda}_k - \hat{\sigma}^2) \hat{u}_k \hat{u}_k^T \end{aligned} \tag{3.2}$$

to estimate r , σ^2 and Θ^* , respectively. Here, μ is a regularization parameter specified by (3.4) in Theorem 3.1 below, and is based on the bounds on $\|\hat{\Sigma} - \Sigma\|_2$ established earlier. Then we let

$$\tilde{\Sigma}^e = \hat{\Theta} + I_d \tag{3.3}$$

be the closed-form estimator of Σ . Note that we do not require $\tilde{\Sigma}^e = \hat{\Theta} + \hat{\sigma}^2 I_d$. Such a requirement could be imposed by solving a convex program like (3.13) with the additional constraint that the diagonal elements of Θ are all equal and are between 0 and 1, but in this section we focus on closed-form estimators.

Note that, by the construction of $\hat{\Theta}$ as in (3.2), the estimated nonzero eigenvalues of $\hat{\Theta}$, namely $\hat{\lambda}_k - \hat{\sigma}^2$ for $1 \leq k \leq \hat{r}$, are always positive. Thus, $\hat{\Theta}$ is positive semidefinite. On the other hand, $\hat{\sigma}^2$ may become negative in the pathological case when $\hat{\Sigma}$ is not positive semidefinite. To address this problem, we could impose a large enough lower bound on σ^2 so that $\hat{\sigma}^2 > 0$ with high probability. Alternatively, we could replace $\hat{\Sigma}$ by its positive semidefinite version $\hat{\Sigma}^+$ as constructed in Corollary 2.5 from the very beginning, and avoid the pathological case altogether. With the bound on $\|\hat{\Sigma}^+ - \Sigma\|_2$ established in the same corollary, all our analysis will follow except for some minor changes in absolute constants. For brevity we omit the details of these changes.

The following theorem summarizes the performance of our closed-form estimators.

Theorem 3.1. *Let $0 < \alpha < 1/2$, $C_1 = \pi$ and $C_2 = 3\pi^2/16 < 1.86$. We set the regularization parameter μ as*

$$\mu = 2\{C_1\sqrt{\|\widehat{T}\|_2 f^2(n, d, \alpha) + \frac{1}{4}f^4(n, d, \alpha)} + (\frac{1}{2}C_1 + C_2)f^2(n, d, \alpha)\}, \quad (3.4)$$

and set

$$\bar{\mu} = 2\{C_1\sqrt{\|T\|_2}f(n, d, \alpha) + (C_1 + C_2)f^2(n, d, \alpha)\}. \quad (3.5)$$

Suppose that Θ^* satisfies $0 < r < d$ and $\lambda_r(\Theta^*) \geq 2\bar{\mu}$, and n is large enough such that inequality (2.2) holds. Then, on an event with probability exceeding $1 - 2\alpha$,

$$\widehat{r} = r, \quad (3.6)$$

$$\|\widetilde{\Sigma}^e - \Sigma\|_F^2 \leq \|\widehat{\Theta} - \Theta^*\|_F^2 \leq 2r\bar{\mu}^2, \quad (3.7)$$

$$|\widehat{\sigma}^2 - \sigma^2| \leq \frac{1}{2}\bar{\mu} \quad (3.8)$$

hold simultaneously. If, in addition, the common value of the diagonal elements of Θ^* is upper bounded by $1 - \sqrt{2r}\bar{\mu}^2$, then $\widetilde{\Sigma}^e$ is positive semidefinite on the same event.

Proof. The proof can be found in Section 5. □

We elaborate the results presented in Theorem 3.1. First, the regularization parameter μ , and hence our closed-form estimators are constructed entirely with explicit constants and measurable quantities. In addition, in the regime specified by (2.2), that is, (roughly) when $n\|T\|_2 \gtrsim d \log(2\alpha^{-1}d)$, the rate $2r\bar{\mu}^2 = \mathcal{O}(\|T\|_2 \cdot rd \log(2\alpha^{-1}d)/n)$ in (3.7) is, up to the operator norm factor $\|T\|_2$ and the logarithmic factor $\log(2\alpha^{-1}d)$, proportional to the number of parameters in the model divided by the sample size. Hence, our estimation procedure achieves correct rank identification for the low-rank component Θ^* , and near-optimal recovery rate in terms of Frobenius norm deviation for both Θ^* and the copula correlation matrix Σ , in a fully data-driven manner.

Theorem 3.1 also shows that, under appropriate conditions, if the diagonal elements of Θ^* are sufficiently less than one, then the estimator $\widetilde{\Sigma}^e$ is positive semidefinite with high probability. In any case, if $\widetilde{\Sigma}^e$ is not positive semidefinite, we can employ Theorem 2.4 to obtain from $\widetilde{\Sigma}^e$ a positive semidefinite estimator $\widetilde{\Sigma}^{e+}$ of Σ such that $\|\widetilde{\Sigma}^{e+} - \Sigma\|_F$ is comparable to $\|\widetilde{\Sigma}^e - \Sigma\|_F$. We defer the details of this treatment to Corollary 3.3.

3.2. Analysis of the refined estimator: Preliminaries

We denote

$$r^* = \text{rank}(\Theta^*).$$

Let Θ^* have the eigen-decomposition

$$\Theta^* = U^* \text{diag}^*(\lambda_1(\Theta^*), \dots, \lambda_{r^*}(\Theta^*))U^{*T}.$$

Here, $\lambda_1(\Theta^*) \geq \dots \geq \lambda_{r^*}(\Theta^*)$ are the positive eigenvalues of Θ^* in descending order, and

$$U^* = (u^1, \dots, u^{r^*})$$

is the $d \times r^*$ matrix of the orthonormal eigenvectors of Θ^* , with the eigenvector u^i corresponding to the eigenvalue $\lambda_i(\Theta^*)$.

Furthermore, for all r with $0 \leq r \leq r^*$, we let

$$U_r^* = (u^1, \dots, u^r) \tag{3.9}$$

be the $d \times r$ truncated matrix of orthonormal eigenvectors of Θ^* , let

$$\gamma_r = \|U_r^* U_r^{*T}\|_\infty, \tag{3.10}$$

and let Θ_r^* be the best rank- r approximation to Θ^* in the Frobenius norm, that is, $\Theta_r^* = \arg \min_{\Theta \in \mathbb{R}^{d \times d}, \text{rank}(\Theta)=r} \|\Theta - \Theta^*\|_F$. We note that γ_r is nondecreasing in r on $0 \leq r \leq r^*$, and $\gamma_{r^*} \leq 1$. In addition, by Schmidt's approximation theorem [39] or the Eckart–Young theorem [12], for $0 \leq r \leq r^*$, we have

$$\Theta_r^* = U_r^* \text{diag}^*(\lambda_1(\Theta^*), \dots, \lambda_r(\Theta^*)) U_r^{*T}, \tag{3.11}$$

and $\|\Theta_r^* - \Theta^*\|_F^2 = \sum_{j: r < j \leq r^*} \lambda_j^2(\Theta^*)$.

3.3. Analysis of the refined estimator: Main result

We first observe that in the elliptical copula correlation factor model, alternative to (1.6), we can write the copula correlation matrix Σ as

$$\Sigma = \Theta_o^* + I_d.$$

This motivates us to set our refined estimator $\tilde{\Sigma}$ of Σ to be

$$\tilde{\Sigma} = \tilde{\Theta}_o + I_d. \tag{3.12}$$

Here, $\tilde{\Theta}$ is our estimator of the low-rank component Θ^* , and is obtained as the solution to a convex program:

$$\tilde{\Theta} = \arg \min_{\Theta \in \mathbb{R}^{d \times d}} \left\{ \frac{1}{2} \|\Theta_o - \hat{\Sigma}_o\|_F^2 + \mu \|\Theta\|_* \right\}. \tag{3.13}$$

(By its optimality, $\tilde{\Theta}$ must be symmetric, though this particular property is not used in our subsequent analysis.) In (3.13), μ is a regularization parameter chosen according to (3.15) in Theorem 3.2 below, and is based on the bounds on $\|\hat{\Sigma} - \Sigma\|_2$ established earlier.

We now elaborate the construction of the refined estimator. Note that:

1. In the factor model, the off-diagonal elements of Σ and Θ^* agree, so the off-diagonal elements of $\hat{\Sigma}$ are natural estimators of the corresponding elements of Θ^* ;

2. The plug-in estimator $\widehat{\Sigma}$, similar to the target copula correlation matrix Σ , has all its diagonal elements equal to one irrespective of the low-rank component Θ^* . As a consequence, we critically lack estimators for the diagonal elements of Θ^* .

Because of these observations, when constructing the estimator $\widetilde{\Theta}$ of Θ^* through the convex program (3.13), we minimize the Frobenius norm for only the off-diagonal elements of the deviation between $\widehat{\Sigma}$ and the estimator of Θ^* subject to a penalty. The penalty is the nuclear norm of the estimator of Θ^* scaled by the regularization parameter μ , and is implemented to encourage the estimator of Θ^* to be appropriately low-rank while keeping (3.13) convex [15]. Then, when constructing the refined estimator $\widetilde{\Sigma}$ of Σ from the estimator $\widetilde{\Theta}$ of Θ^* through (3.12), we explicitly set all the diagonal elements of $\widetilde{\Sigma}$ to one. It is clear that any bound on $\widetilde{\Sigma} - \Sigma$ also acts as a bound on the off-diagonal elements of $\widetilde{\Theta} - \Theta^*$ and vice versa. We bound the diagonal elements of $\widetilde{\Theta} - \Theta^*$ in Appendix C.

We briefly contrast our refined estimator $\widetilde{\Sigma}$, which is tailor-made for our special setting of the elliptical copula correlation factor model, to some of the existing estimation procedures in related but different contexts.

1. Our setting is an extension of the low-rank matrix approximation problem [31,34,37]. In particular, [31] studies the estimation of Θ^* that is a covariance matrix⁴ with low effective rank, with the added complication that the observations X^1, \dots, X^n are masked at *random* coordinates. [31] constructs an unbiased initial estimator $\widehat{\Theta}$ of Θ^* , and further obtains a refined estimator $\widetilde{\Theta}$ as the solution of a convex program that is identical to (3.13) but with the term $\|\Theta_o - \widehat{\Sigma}_o\|_F^2$ replaced by $\|\Theta - \widehat{\Theta}\|_F^2$, which is a sum over *all* entries of the matrix $\Theta - \widehat{\Theta}$.

Contrary to the setting of [31], Σ in the factor model (1.6) typically has neither low effective rank nor low rank: because $\text{tr}(\Sigma) = d$, the effective rank of Σ is $r_e(\Sigma) = d/\|\Sigma\|_2$, which is large unless $\|\Sigma\|_2$ becomes comparable to d ; in addition, because Θ^* is positive semidefinite, if the diagonal elements of V^* are all strictly positive, then $\Sigma = \Theta^* + V^*$ has full rank. Hence, a naive application of the method of [31] to our setting amounts to seeking a low-rank approximation to a matrix that is in fact not low-rank. In contrast, our program (3.13) seeks to estimate the genuine low-rank or nearly low-rank component Θ^* of Σ , even though this choice leads to technical challenges in our proof as compared to [31].

2. By the observations we made earlier, our problem can be rephrased as follows: Estimate the off-diagonal elements of Θ^* given only their noisy observations, taking advantage of the fact that Θ^* is low-rank or nearly low-rank. Hence, as mentioned in Section 1.2, our problem is a variant of the matrix completion problem, in particular the version in which a matrix Σ (not necessarily a correlation matrix) admits a decomposition into the sum of a low-rank component Θ^* and a sparse component S^* with a general sparsity pattern (i.e., the locations of the nonzero entries of the sparse component are unknown but fixed), and the goal is to estimate Σ based on its noisy observation $\widehat{\Sigma}$ [1,9,10,20,32,48]. In particular, [9,20] let $\widetilde{\Theta}$, the estimator of Θ^* , and \widetilde{S} , the estimator of S^* , be the solution of

$$(\widetilde{\Theta}, \widetilde{S}) = \arg \min_{\Theta, S \in \mathbb{R}^{d \times d}} \left\{ \frac{1}{2} \|\Theta + S - \widehat{\Sigma}\|_F^2 + \mu \|\Theta\|_* + \lambda \|S\|_1 \right\}. \tag{3.14}$$

⁴For this paragraph only, we use Θ^* to denote the covariance matrix, because in the setting of [31] it is the covariance matrix itself that has low effective rank.

This scenario is the closest to our setting. However, even though V^* in the factor model is indeed a sparse matrix, and thus one could apply (3.14) to our setting, such an approach would not be optimal because it obviously takes no advantage of our knowledge of the sparsity pattern of V^* , namely the diagonal pattern. For instance, [9,20] require nontrivial specification of an additional regularization parameter $\lambda = \lambda(\mu)$ for the element-wise ℓ_1 penalty of the sparse component. Because (3.13) and (3.14) are distinct programs, it is also not possible to infer the properties of our refined estimator $\tilde{\Sigma}$ directly from the results of [9,20].

3. Finally, the low-rank and diagonal matrix decomposition problem in the *noiseless* setting is treated in [38]. These authors employ a semidefinite program, the minimum trace factor analysis (MTFA), to minimize the trace of the low-rank component (subject to the constraint that the sum of the low-rank component and the diagonal component agrees with the given matrix to be decomposed). The optimality condition from semidefinite programming then gives fairly simple conditions for the MTFA to exactly recover the decomposition.

We adopt the *primal-dual certificate* approach advocated by [20,48]⁵ to analyze (3.13). Our oracle inequality for the refined estimator $\tilde{\Sigma}$ is collected in the following theorem.

Theorem 3.2. *Let $0 < \alpha < 1/2$, $C_1 = \pi$, $C_2 = 3\pi^2/16 < 1.86$, and $C = 6$. We set the regularization parameter μ as*

$$\mu = C \left\{ C_1 \sqrt{\|\hat{T}\|_2 f^2(n, d, \alpha) + \frac{1}{4} f^4(n, d, \alpha) + \left(\frac{1}{2}C_1 + C_2\right) f^2(n, d, \alpha)} \right\}, \tag{3.15}$$

and set

$$\bar{\mu} = C \left\{ C_1 \max[\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)] + (C_1 + C_2) f^2(n, d, \alpha) \right\}. \tag{3.16}$$

Recall γ_r as defined in (3.10). We set

$$R = \max\{r: 0 \leq r \leq r^*, \gamma_r \leq 1/9\}. \tag{3.17}$$

Then, with probability exceeding $1 - 2\alpha$, the refined estimator $\tilde{\Sigma}$, as introduced in (3.12), of Σ satisfies

$$\|\tilde{\Sigma} - \Sigma\|_F^2 \leq \min_{0 \leq r \leq R} \left\{ \sum_{j: r < j \leq r^*} \lambda_j^2(\Theta^*) + 8r\bar{\mu}^2 \right\}. \tag{3.18}$$

Remark. Theorem 3.2 is a specific instance of Corollary 6.8 which is a more general result; in particular the constant $C = 6$ in (3.15) and (3.16) and the upper bound $1/9$ on γ_r in (3.17) are chosen for ease of presentation but are not specifically optimized. For instance, we could specify a smaller C at the expense of a more stringent upper bound on γ_r .

⁵Through delicate analysis, [9] (which builds upon their earlier work [10] in the noiseless setting) guarantees optimal convergence rate in terms of the operator norm, as well as consistent rank recovery, for the estimator $\tilde{\Theta}$ of the low-rank component Θ^* . On the other hand, their analysis requires that the minimum nonzero singular value of the low-rank component Θ^* satisfies a nontrivial lower bound, and hence at this stage is not particularly well suited to study the case where the low-rank requirement only holds approximately.

Proof of Theorem 3.2. The proof can be found in Section 6. □

We elaborate the results presented in Theorem 3.2.

The oracle inequality (3.18) in fact represents the minimum of a collection of upper bounds, and the minimum is taken over all r that satisfies $\gamma_r = \|U_r^* U_r^{*T}\|_\infty \leq 1/9$, a range specified by (3.17). Thus, for the oracle inequality (3.18) to be as tight as possible, we should ideally have a large range of r such that $\gamma_r \leq 1/9$. We discuss two concrete examples in which this condition is satisfied:

1. If for some given r , the entries of u^i , $1 \leq i \leq r$ are all bounded by c/\sqrt{d} for some constant $c \geq 1$, then $\gamma_r \leq c^2 r/d$;

2. Next, we consider the *random orthogonal model* as in [8]. The first result of their Lemma 2.2 shows that, if u^i , $1 \leq i \leq r$ are sampled uniformly at random among all families of r orthonormal vectors independently of each other, then there exist constants C and c such that $\gamma_r \leq C \max\{r, \log(d)\}/d$ with probability at least $1 - cd^{-3} \log d$.

In both cases, $\gamma_r \leq 1/9$ is satisfied for all r 's that are small compared to d (in the second case when d is large enough and with high probability to be precise).

The estimation procedure (3.13) is fully data-driven; in particular, the penalty term in (3.13) is scaled by a regularization parameter μ specified by (3.15) with explicit constants and measurable quantities. In addition, procedure (3.13) automatically balances the approximation error with the estimation error as if it knows the right model in advance to arrive at the oracle inequality (3.18) with near-optimal recovery rate in terms of Frobenius norm deviation. Specifically,

1. The primal-dual certificate approach yields an approximation error term, that is, the first term in the curly bracket on the right-hand side of (3.18), with leading multiplicative constant one. Such a feature has become increasingly common with the results obtained through convex optimization with nuclear norm penalty [25,31];

2. Meanwhile, the estimation error term, that is, the second term in the curly bracket on the right-hand side of (3.18), achieves a rate $8r\bar{\mu}^2 = \mathcal{O}(\|T\|_2 \cdot rd \log(2\alpha^{-1}d)/n)$ with probability exceeding $1 - 2\alpha$ if we focus on the regime specified by (2.2), that is, (roughly) when $n\|T\|_2 \gtrsim d \log(2\alpha^{-1}d)$. Again, this rate is, up to the operator norm factor $\|T\|_2$ and the logarithmic factor $\log(2\alpha^{-1}d)$, proportional to the number of parameters in the model divided by the sample size.⁶

⁶Again by the lower bound argument presented in the proof of [31], Theorem 2, the rate of the estimation error term in (3.18) is optimal up to the operator norm factor and the log factor. We note that the lower bounds (and in particular the one for Frobenius norm deviation) established by [31], Theorem 2, contain explicit dependence on the operator norm $\|\Sigma\|_2$ of the target covariance matrix Σ in the form of a multiplicative factor. However, a closer inspection of the proof of [31], Theorem 2, reveals that this particular $\|\Sigma\|_2$ is in fact restricted to be at most two times the maximum of the diagonal elements of Σ , and thus in our case can at most be two because Σ is a correlation matrix. This restriction is not ideal because $\|\Sigma\|_2$ in general can be as large as d . In our opinion, it remains to be seen how a proper dependence on operator norm can be obtained in lower bound for Frobenius norm deviation under our setting of correlation matrix estimation. From another angle, we have shown in the proof of Corollary 2.5 that the plug-in estimator $\hat{\Sigma}$ achieves $\|\hat{\Sigma} - \Sigma\|_\infty = \mathcal{O}(\sqrt{\log(2\alpha^{-1}d)/n})$ (with probability at least $1 - \frac{1}{4}\alpha^2$); thus $\|\hat{\Sigma} - \Sigma\|_F^2 = \mathcal{O}(d^2 \cdot \log(2\alpha^{-1}d)/n)$ (with the same probability). This rate is slower than $r\bar{\mu}^2$ so long as $r\|T\|_2 \lesssim d$. Therefore, the presence of $\|T\|_2$ in (3.7) and (3.18) entails an upper bound on the rank of the low-rank component Θ^* below which the refined estimator and the closed form estimator in their respective contexts are preferable to the plug-in estimator $\hat{\Sigma}$ in terms of Frobenius norm deviation.

Finally, if the diagonal elements of the deviation $\tilde{\Theta} - \Theta^*$ can be appropriately bounded, for instance, through Theorem C.2 in Appendix C, and if the diagonal elements of Θ^* are sufficiently smaller than one, then the estimator $\tilde{\Sigma}$ is positive semidefinite. Because the argument is similar to the proof of the last statement of Theorem 3.1, we omit its details. In any case, if $\tilde{\Sigma}$ is not positive semidefinite, we can employ Theorem 2.4 to obtain from $\tilde{\Sigma}$ a positive semidefinite estimator $\tilde{\Sigma}^+$ of Σ such that $\|\tilde{\Sigma}^+ - \Sigma\|_F$ is comparable to $\|\tilde{\Sigma} - \Sigma\|_F$, as Corollary 3.3 demonstrates.

Corollary 3.3. *In (2.9), we let the generic matrix norm $\|\cdot\|$ be replaced by the Frobenius norm $\|\cdot\|_F$, and let $\mathcal{F} = \mathcal{S}_+^d$. In addition, in the context of the elementary factor copula model, we let the generic estimator $\widehat{\Sigma}^{\text{generic}}$ be replaced by the closed-form estimator $\tilde{\Sigma}^e$, and the solution $\widehat{\Sigma}^{\text{generic}+}$ be replaced by $\tilde{\Sigma}^{e+}$, while in the context of the (general) elliptical copula correlation factor model, we let the generic estimator $\widehat{\Sigma}^{\text{generic}}$ be replaced by the refined estimator $\tilde{\Sigma}$, and the solution $\widehat{\Sigma}^{\text{generic}+}$ be replaced by $\tilde{\Sigma}^+$. Then $\tilde{\Sigma}^{e+}$ and $\tilde{\Sigma}^+$ satisfy*

$$\|\tilde{\Sigma}^{e+} - \Sigma\|_F \leq 2\|\tilde{\Sigma}^e - \Sigma\|_F, \quad \|\tilde{\Sigma}^+ - \Sigma\|_F \leq 2\|\tilde{\Sigma} - \Sigma\|_F. \tag{3.19}$$

We recall that $\|\tilde{\Sigma}^e - \Sigma\|_F$ and $\|\tilde{\Sigma} - \Sigma\|_F$ are bounded as in Theorems 3.1 and 3.2, respectively.

Remark. We refer the readers to [36] and the references therein for the computational aspect of (2.9) in this context of Frobenius norm minimization.

Proof of Corollary 3.3. With the choice $\mathcal{F} = \mathcal{S}_+^d$, we clearly have $\Sigma \in \mathcal{F}$. Then (3.19) follows straightforwardly from Theorem 2.4. \square

For both Corollaries 2.5 and 3.3, we have obtained positive semidefinite, rather than strictly positive definite, versions of the existing estimators. To obtain strictly positive definite estimators, we could replace the existing feasible regions \mathcal{F} in Corollaries 2.5 and 3.3 by an intersection of \mathcal{F} and the convex set $\{\Sigma' \in \mathbb{R}^{d \times d} : \lambda_{\min}(\Sigma') \geq \varepsilon\}$ for some $\varepsilon > 0$. Then the resulting estimator from (2.9) will be positive definite, with the smallest eigenvalue lower bounded by ε . If in addition the copula correlation matrix Σ satisfies $\lambda_{\min}(\Sigma) \geq \varepsilon$, the conclusions of Corollaries 2.5 and 3.3 will continue to hold.

4. Proofs for Section 2

4.1. Proof of Theorem 2.1

The proof of Theorem 2.1 is further divided into two stages. In Section 4.1.1, we prove inequality (2.1a); in Section 4.1.2, we prove the data-driven bound, inequality (2.1b), and its performance guarantee, inequality (2.1c).

4.1.1. Proof of inequality (2.1a)

We wish to apply a Bernstein-type inequality, specifically [42], Theorem 6.6.1, to bound the tail probability $\mathbb{P}\{\|\widehat{T} - T\|_2 \geq t\}$. We note that this theorem on bounding the tail probability of the

maximum eigenvalue of a sum of random matrices requires that the summands be independent. Clearly, the matrix U -statistic $\widehat{T} - T$ does not satisfy this condition. On the other hand, this theorem relies on the Chernoff transform technique to convert the tail probability into an expectation of a convex function of $\widehat{T} - T$. A technique by Hoeffding [18] then allows us to convert the problem of bounding $\|\widehat{T} - T\|_2$ into a problem involving a sum of independent random matrices.

Proposition 4.1. *We define*

$$\widetilde{T} = \frac{2}{n} \sum_{i=1}^{n/2} \widetilde{T}^i \tag{4.1}$$

with

$$\widetilde{T}^i = \text{sgn}(X^{2i-1} - X^{2i}) \text{sgn}(X^{2i-1} - X^{2i})^T. \tag{4.2}$$

Then the tail probability $\mathbb{P}\{\|\widehat{T} - T\|_2 \geq t\}$ satisfies

$$\mathbb{P}\{\|\widehat{T} - T\|_2 \geq t\} \leq \inf_{\theta > 0} \{e^{-\theta t} \cdot \mathbb{E}[\text{tr} e^{\theta(\widehat{T}-T)}]\} + \inf_{\theta > 0} \{e^{-\theta t} \cdot \mathbb{E}[\text{tr} e^{\theta(T-\widetilde{T})}]\}.$$

Proof. First, note that, because $\widehat{T} - T$ is symmetric, we have

$$\|\widehat{T} - T\|_2 = \max\{\lambda_{\max}(\widehat{T} - T), -\lambda_{\min}(\widehat{T} - T)\} = \max\{\lambda_{\max}(\widehat{T} - T), \lambda_{\max}(T - \widehat{T})\}.$$

Hence,

$$\begin{aligned} \mathbb{P}\{\|\widehat{T} - T\|_2 \geq t\} &= \mathbb{P}\{\{\lambda_{\max}(\widehat{T} - T) \geq t\} \cup \{\lambda_{\max}(T - \widehat{T}) \geq t\}\} \\ &\leq \mathbb{P}\{\lambda_{\max}(\widehat{T} - T) \geq t\} + \mathbb{P}\{\lambda_{\max}(T - \widehat{T}) \geq t\}. \end{aligned} \tag{4.3}$$

Next we bound the first term on the right-hand side of inequality (4.3), that is, $\mathbb{P}\{\lambda_{\max}(\widehat{T} - T) \geq t\}$. Applying the Chernoff transform technique (e.g., [42], Proposition 3.2.1), we have

$$\mathbb{P}\{\lambda_{\max}(\widehat{T} - T) \geq t\} \leq \inf_{\theta > 0} \{e^{-\theta t} \cdot \mathbb{E}[\text{tr} e^{\theta(\widehat{T}-T)}]\}. \tag{4.4}$$

Now we introduce the technique of Hoeffding. We note the following facts:

1. We can equivalently write \widehat{T} as

$$\widehat{T} = \frac{1}{n!} \sum_{n,n} V(X^{i_1}, \dots, X^{i_n}). \tag{4.5}$$

Here, the function V is defined as

$$V(X^{i_1}, \dots, X^{i_n}) = \frac{2}{n} \{g(X^{i_1}, X^{i_2}) + g(X^{i_3}, X^{i_4}) + \dots + g(X^{i_{n-1}}, X^{i_n})\},$$

the kernel g is defined as

$$g(X^i, X^j) = \text{sgn}(X^i - X^j) \text{sgn}(X^i - X^j)^T,$$

and the sum $\sum_{n,n}$ is taken over all permutations i_1, i_2, \dots, i_n of the integers $1, 2, \dots, n$.

2. The trace exponential function is convex on the set of Hermitian matrices [35].

Therefore, using first (4.5) and then Jensen’s inequality, we have

$$\begin{aligned} \text{tr} e^{\theta(\widehat{T}-T)} &= \text{tr} \exp \left\{ \sum_{n,n} \frac{1}{n!} \theta [V(X^{i_1}, \dots, X^{i_n}) - T] \right\} \\ &\leq \sum_{n,n} \frac{1}{n!} \text{tr} \exp \{ \theta [V(X^{i_1}, \dots, X^{i_n}) - T] \}. \end{aligned} \tag{4.6}$$

Then, plugging inequality (4.6) into inequality (4.4), we have

$$\begin{aligned} \mathbb{P} \{ \lambda_{\max}(\widehat{T} - T) \geq t \} &\leq \inf_{\theta > 0} \left\{ e^{-\theta t} \cdot \mathbb{E} \left[\sum_{n,n} \frac{1}{n!} \text{tr} e^{\theta [V(X^{i_1}, \dots, X^{i_n}) - T]} \right] \right\} \\ &= \inf_{\theta > 0} \{ e^{-\theta t} \cdot \mathbb{E} [\text{tr} e^{\theta [V(X^1, X^2, \dots, X^n) - T]}] \} \\ &= \inf_{\theta > 0} \{ e^{-\theta t} \cdot \mathbb{E} [\text{tr} e^{\theta(\widetilde{T}-T)}] \}. \end{aligned}$$

The second term on the right-hand side of inequality (4.3) can be similarly bounded. The conclusion of the proposition then follows. \square

In Proposition 4.1, the argument of the trace exponential function is proportional to

$$\widetilde{T} - T = \sum_{i=1}^{n/2} \frac{2}{n} (\widetilde{T}^i - T),$$

with now independent summands $2n^{-1}(\widetilde{T}^i - T)$, $1 \leq i \leq n/2$, which are also symmetric. Therefore, we can proceed as in the proof of [42], Theorem 6.6.1, to bound $\mathbb{E}[\text{tr} e^{\theta(\widetilde{T}-T)}]$ and $\mathbb{E}[\text{tr} e^{\theta(T-\widetilde{T})}]$. We calculate the quantities necessary for applying the proof. First, (for any i) we clearly have $\mathbb{E}[\widetilde{T}^i - T] = 0$. Next, by the representation of \widetilde{T}^i as in (4.2), we conclude that the only nonzero eigenvalue of \widetilde{T}^i is d which corresponds to the eigenvector $\text{sgn}(X^{2i-1} - X^{2i})$; thus, $\lambda_{\max}(\widetilde{T}^i) = d$. This, together with Weyl’s inequality and the facts that T is positive semidefinite and $\|T\|_2 \leq d \cdot \|T\|_\infty \leq d$, imply that

$$\lambda_{\max}(\widetilde{T}^i - T) \leq \lambda_{\max}(\widetilde{T}^i) = d, \tag{4.7a}$$

$$\lambda_{\max}(T - \widetilde{T}^i) \leq \lambda_{\max}(T) \leq d. \tag{4.7b}$$

Finally, we calculate

$$\sigma^2 = \left\| \sum_{i=1}^{n/2} \mathbb{E} \left\{ \left[\frac{2}{n} (\tilde{T}^i - T) \right]^2 \right\} \right\|_2,$$

the matrix variance statistic of the sum as defined in [42], Theorem 6.6.1. Note that

$$\begin{aligned} (\tilde{T}^i)^2 &= \text{sgn}(X^{2i-1} - X^{2i}) \text{sgn}(X^{2i-1} - X^{2i})^T \text{sgn}(X^{2i-1} - X^{2i}) \text{sgn}(X^{2i-1} - X^{2i})^T \\ &= \text{sgn}(X^{2i-1} - X^{2i}) [\text{sgn}(X^{2i-1} - X^{2i})^T \text{sgn}(X^{2i-1} - X^{2i})] \text{sgn}(X^{2i-1} - X^{2i})^T \\ &= d \cdot \text{sgn}(X^{2i-1} - X^{2i}) \text{sgn}(X^{2i-1} - X^{2i})^T \\ &= d \cdot \tilde{T}^i. \end{aligned}$$

Then

$$\left(\frac{n}{2}\right)^2 \sigma^2 = \left\| \sum_{i=1}^{n/2} \mathbb{E} [d \cdot \tilde{T}^i - T^2] \right\|_2 = \frac{n}{2} \|d \cdot T - T^2\|_2 \leq \frac{n}{2} d \|T\|_2. \tag{4.8}$$

Hence, by Proposition 4.1 and the proof of [42], inequality (6.6.3) in Theorem 6.6.1, as well as (4.7a), (4.7b) and (4.8), we obtain the matrix Bernstein inequality

$$\begin{aligned} \mathbb{P}(\|\hat{T} - T\|_2 \geq t) &\leq 2d \cdot \exp\left(-\frac{nt^2}{4d\|T\|_2 + 4dt/3}\right) \\ &\leq 2d \cdot \max\left\{\exp\left(-\frac{3}{16} \frac{nt^2}{d\|T\|_2}\right), \exp\left(-\frac{3}{16} \frac{nt}{d}\right)\right\}. \end{aligned} \tag{4.9}$$

(By Proposition 4.1 and the proof of [41], Theorem 6.1, we can also obtain the tighter matrix Bennett inequality.) Finally, setting the right-hand side of inequality (4.9) to α and solving for t yields that inequality (2.1a) holds with probability at least $1 - \alpha$.

4.1.2. Proof of inequalities (2.1b) and (2.1c)

We abbreviate $f(n, d, \alpha)$ by f , $\|T\|_2$ by t , $\|\hat{T}\|_2$ by \hat{t} , and $\|\hat{T} - T\|_2$ by δ . We have already established that we have an event with probability at least $1 - \alpha$ on which inequality (2.1a), that is, $\delta < \max\{f\sqrt{t}, f^2\}$, holds, and we concentrate on this event.

We proceed to prove inequality (2.1b), which states

$$\max\{f\sqrt{t}, f^2\} \leq \sqrt{\hat{t}f^2 + \left(\frac{1}{2}f^2\right)^2} + \frac{1}{2}f^2. \tag{4.10}$$

Now, if $f\sqrt{t} \leq f^2$ and so $\max\{f\sqrt{t}, f^2\} = f^2$, then inequality (4.10) clearly holds. Thus, we focus on the case $f\sqrt{t} > f^2$. In this case, by inequality (2.1a), we must have

$$\delta < f\sqrt{t}. \tag{4.11}$$

By the triangle inequality,

$$f\sqrt{t} \leq f\sqrt{\delta + \hat{t}}. \tag{4.12}$$

Then, from inequalities (4.11) and (4.12) we deduce

$$f\sqrt{t} < f\sqrt{f\sqrt{t} + \hat{t}}. \tag{4.13}$$

Squaring both sides of inequality (4.13) yields $tf^2 < f^3\sqrt{t} + \hat{t}f^2$, or equivalently

$$(f\sqrt{t} - \frac{1}{2}f^2)^2 < \hat{t}f^2 + (\frac{1}{2}f^2)^2. \tag{4.14}$$

Because in the current case $f\sqrt{t} > f^2 > \frac{1}{2}f^2$, inequality (4.14) implies

$$f\sqrt{t} < \sqrt{\hat{t}f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2,$$

which, together with $f\sqrt{t} > f^2$, again implies inequality (4.10). Hence, we have proved inequality (2.1b).

Next, we prove inequality (2.1c). By the triangle inequality,

$$\sqrt{\hat{t}f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 \leq \sqrt{tf^2 + \delta f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2. \tag{4.15}$$

First, assume that $\delta < f\sqrt{t}$. Then, from inequality (4.15) we deduce

$$\begin{aligned} \sqrt{\hat{t}f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 &< \sqrt{tf^2 + f^3\sqrt{t} + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 \\ &= (f\sqrt{t} + \frac{1}{2}f^2) + \frac{1}{2}f^2. \end{aligned} \tag{4.16}$$

Next, suppose instead $\delta \geq f\sqrt{t}$, so by inequality (2.1a) we must have $f\sqrt{t} \leq \delta < f^2$. Then, from inequality (4.15) we deduce

$$\sqrt{\hat{t}f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 < \sqrt{f^4 + f^4 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 = \frac{3}{2}f^2 + \frac{1}{2}f^2. \tag{4.17}$$

Both inequalities (4.16) and (4.17) further imply that

$$\sqrt{\hat{t}f^2 + (\frac{1}{2}f^2)^2} + \frac{1}{2}f^2 < \max\{f\sqrt{t}, f^2\} + f^2,$$

which is just inequality (2.1c).

4.2. Proof of Theorem 2.2

The proof of Theorem 2.2 will be established through the following three lemmas. Recall that we use \circ to denote the Hadamard product.

Lemma 4.2. *We have*

$$\|\widehat{\Sigma}' - \Sigma\|_2 \leq \frac{\pi}{2} \cdot \left\| \cos\left(\frac{\pi}{2}T\right) \circ (\widehat{T}' - T) \right\|_2 + \frac{\pi^2}{8} \cdot \left\| \sin\left(\frac{\pi}{2}\overline{T}\right) \circ (\widehat{T}' - T) \circ (\widehat{T}' - T) \right\|_2.$$

Here, \overline{T} is a symmetric, random matrix such that each entry $[\overline{T}]_{k\ell}$ is a random number on the closed interval between $[T]_{k\ell}$ and $[\widehat{T}']_{k\ell}$.

Proof. By Taylor’s theorem, we have

$$\begin{aligned} \widehat{\Sigma}' - \Sigma &= \sin\left(\frac{\pi}{2}\widehat{T}'\right) - \sin\left(\frac{\pi}{2}T\right) \\ &= \cos\left(\frac{\pi}{2}T\right) \circ \frac{\pi}{2}(\widehat{T}' - T) - \frac{1}{2} \sin\left(\frac{\pi}{2}\overline{T}\right) \circ \frac{\pi}{2}(\widehat{T}' - T) \circ \frac{\pi}{2}(\widehat{T}' - T), \end{aligned} \tag{4.18}$$

for some matrix \overline{T} as specified in the theorem. Next, applying the operator norm on both sides of equation (4.18) and then using the triangle inequality on the right-hand side yields the lemma. \square

Hence, it suffices to establish appropriate bounds separately for a first-order term, $\|\cos(\frac{\pi}{2}T) \circ (\widehat{T}' - T)\|_2$, and a second-order term, $\|\sin(\frac{\pi}{2}\overline{T}) \circ (\widehat{T}' - T) \circ (\widehat{T}' - T)\|_2$.

Lemma 4.3. *For the first-order term, we have*

$$\left\| \cos\left(\frac{\pi}{2}T\right) \circ (\widehat{T}' - T) \right\|_2 \leq 2\|\widehat{T}' - T\|_2.$$

Proof. Recall that $\sin(\frac{\pi}{2}T) = \Sigma$. Then, with J_d denoting a $d \times d$ matrix with all entries identically equal to one, and the square root function acting component-wise, we have

$$\cos\left(\frac{\pi}{2}T\right) = \sqrt{J_d - \sin\left(\frac{\pi}{2}T\right) \circ \sin\left(\frac{\pi}{2}T\right)} = \sqrt{J_d - \Sigma \circ \Sigma}. \tag{4.19}$$

Next, using the generalized binomial formula

$$(1 + x)^\alpha = \sum_{k=0}^{\infty} \binom{\alpha}{k} x^k$$

on equation (4.19) with $\alpha = \frac{1}{2}$ and x being the components of $-\Sigma \circ \Sigma$ (so the sum converges, in fact absolutely, since $\alpha > 0$ and $\|\Sigma \circ \Sigma\|_\infty \leq 1$), we have

$$\cos\left(\frac{\pi}{2}T\right) = \sum_{k=0}^{\infty} \binom{1/2}{k} (-1)^k \Sigma \circ_{2k} \Sigma.$$

Here, by $\Sigma \circ_l \Sigma$ we mean the Hadamard product of l Σ 's, that is, $\Sigma \circ \dots \circ \Sigma$ with a total of l terms. Hence,

$$\begin{aligned} \left\| \cos\left(\frac{\pi}{2}T\right) \circ (\widehat{T}' - T) \right\|_2 &= \left\| \left[\sum_{k=0}^{\infty} \binom{1/2}{k} (-1)^k \Sigma \circ_{2k} \Sigma \right] \circ (\widehat{T}' - T) \right\|_2 \\ &\leq \sum_{k=0}^{\infty} \left| \binom{1/2}{k} \right| \cdot \|(\Sigma \circ_{2k} \Sigma) \circ (\widehat{T}' - T)\|_2. \end{aligned} \tag{4.20}$$

Because Σ is positive semidefinite (since it is a correlation matrix), by the Schur product theorem, $\Sigma \circ_{2k} \Sigma$ is positive semidefinite for all k ; moreover, $\Sigma \circ_{2k} \Sigma$'s all have diagonal elements identically equal to one. Then, by [19], Theorem 5.5.18, we have, for all k ,

$$\|(\Sigma \circ_{2k} \Sigma) \circ (\widehat{T}' - T)\|_2 \leq \|\widehat{T}' - T\|_2. \tag{4.21}$$

Plugging (4.21) into (4.20) and then using the fact that $\sum_{k=0}^{\infty} \binom{1/2}{k} = 2$ yield

$$\left\| \cos\left(\frac{\pi}{2}T\right) \circ (\widehat{T}' - T) \right\|_2 \leq \left[\sum_{k=0}^{\infty} \left| \binom{1/2}{k} \right| \right] \cdot \|\widehat{T}' - T\|_2 = 2\|\widehat{T}' - T\|_2, \tag{4.22}$$

which is the conclusion of the lemma. □

Lemma 4.4. *For the second-order term, we have*

$$\left\| \sin\left(\frac{\pi}{2}\overline{T}\right) \circ (\widehat{T}' - T) \circ (\widehat{T}' - T) \right\|_2 \leq \|\widehat{T}' - T\|_2^2. \tag{4.23}$$

Alternatively, for the particular case $\widehat{T}' = \widehat{T}$, we have, with probability at least $1 - \frac{1}{4}\alpha^2$,

$$\left\| \sin\left(\frac{\pi}{2}\overline{T}\right) \circ (\widehat{T} - T) \circ (\widehat{T} - T) \right\|_2 \leq 8 \cdot \frac{d \cdot \log(2\alpha^{-1}d)}{n}. \tag{4.24}$$

Proof. First, we observe a simple fact: for two matrices $M, N \in \mathbb{R}^{k \times \ell}$ (for arbitrary k, ℓ), if $|[M]_{ij}| \leq [N]_{ij}$ for all $1 \leq i \leq k, 1 \leq j \leq \ell$, then $\|M\|_2 \leq \|N\|_2$.

To see this, we fix an arbitrary vector $u = (u_1, \dots, u_\ell)^T \in \mathbb{R}^\ell$ with $\|u\| = 1$, with $\|\cdot\|$ being the Euclidean norm for vectors. Let $\tilde{u} = (\tilde{u}_1, \dots, \tilde{u}_\ell)^T \in \mathbb{R}^\ell$ be the vector such that $\tilde{u}_j = |u_j|$ for $j = 1, \dots, \ell$, that is, each component of \tilde{u} is the absolute value of the corresponding component of u . Clearly, $\|\tilde{u}\| = 1$ as well. Then we have, for all $1 \leq i \leq k$,

$$|[Mu]_i| = \left| \sum_{j=1}^{\ell} [M]_{ij} u_j \right| \leq \sum_{j=1}^{\ell} |[M]_{ij}| |u_j| \leq \sum_{j=1}^{\ell} [N]_{ij} \tilde{u}_j = |[N\tilde{u}]_i|.$$

Here, $[Mu]_i$ and $[N\tilde{u}]_i$ are the i th component of the vectors Mu and $N\tilde{u}$, respectively. Hence, clearly, $\|Mu\| \leq \|N\tilde{u}\|$, which further implies that

$$\sup\{\|Mu\|: \|u\| = 1\} \leq \sup\{\|Nu\|: \|u\| = 1\},$$

and we conclude that $\|M\|_2 \leq \|N\|_2$.

Now, it is easy to see that

$$\left| \left[\sin\left(\frac{\pi}{2}\bar{T}\right) \circ (\hat{T}' - T) \circ (\hat{T}' - T) \right]_{ij} \right| \leq [(\hat{T}' - T) \circ (\hat{T}' - T)]_{ij} \quad \forall 1 \leq i, j \leq d.$$

Hence, by the preceding observation, we have

$$\left\| \sin\left(\frac{\pi}{2}\bar{T}\right) \circ (\hat{T}' - T) \circ (\hat{T}' - T) \right\|_2 \leq \|(\hat{T}' - T) \circ (\hat{T}' - T)\|_2. \tag{4.25}$$

By [19], Theorem 5.5.1, we further have

$$\|(\hat{T}' - T) \circ (\hat{T}' - T)\|_2 \leq \|\hat{T}' - T\|_2^2. \tag{4.26}$$

Then inequality (4.23) follows from inequalities (4.25) and (4.26).

Next, we prove the second half of the lemma. We have

$$\left\| \sin\left(\frac{\pi}{2}\bar{T}\right) \circ (\hat{T} - T) \circ (\hat{T} - T) \right\|_2 \leq \|(\hat{T} - T) \circ (\hat{T} - T)\|_2 \leq d\|\hat{T} - T\|_\infty^2. \tag{4.27}$$

Here, the first inequality follows by inequality (4.25) with the choice $\hat{T}' = \hat{T}$, and the second inequality follows by the bound that $\|M \circ M\|_2 \leq d\|M \circ M\|_\infty = d\|M\|_\infty^2$ for arbitrary $M \in \mathbb{R}^{d \times d}$. By Hoeffding's inequality for the scalar U -statistic [18],

$$\mathbb{P}(|\hat{T}_{jk} - T_{jk}| \geq t) \leq 2 \exp\left(-\frac{nt^2}{4}\right),$$

and so, by the union bound,

$$\mathbb{P}(\|\hat{T} - T\|_\infty \geq t) \leq d^2 \exp\left(-\frac{nt^2}{4}\right).$$

Thus, there exists an event A with probability at least $1 - \frac{1}{4}\alpha^2$ such that

$$\|\hat{T} - T\|_\infty^2 \leq 4 \cdot \frac{\log(4\alpha^{-2}d^2)}{n} = 8 \cdot \frac{\log(2\alpha^{-1}d)}{n} \tag{4.28}$$

on the event A . Plugging inequality (4.28) into inequality (4.27) yields that inequality (4.24) holds on the same event. This finishes the proof of the lemma. \square

The conclusions of Theorem 2.2 now follow immediately. In particular, inequality (2.5) follows from Lemmas 4.2, 4.3 and inequality (4.23) in Lemma 4.4, while inequality (2.6) follows from Lemmas 4.2 and 4.3 with \widehat{T}' set to \widehat{T} and $\widehat{\Sigma}'$ set to $\widehat{\Sigma}$, and inequality (4.24) in Lemma 4.4, which holds with probability at least $1 - \frac{1}{4}\alpha^2$.

4.3. Proof of Theorem 2.3

We let the arcsin function have the series expansion $\arcsin(x) = \sum_{k=0}^{\infty} g(k)x^k$ for $|x| \leq 1$. The exact form of the $g(k)$'s for all k is not important; we only need $g(0) = 0$, $g(1) = 1$, all the $g(k)$'s are nonnegative, and $\sum_{k=0}^{\infty} g(k) = \pi/2$. With the arcsin function acting component-wise, and with $\Sigma \circ_k \Sigma$ denoting the Hadamard product of k Σ 's, we have

$$T = \frac{2}{\pi} \arcsin(\Sigma) = \frac{2}{\pi} \sum_{k=0}^{\infty} g(k) \Sigma \circ_k \Sigma.$$

Because Σ is positive semidefinite, by the Schur product theorem, $\Sigma \circ_k \Sigma$, and thus $g(k) \Sigma \circ_k \Sigma$, are positive semidefinite for all $k \geq 0$. In addition, T is positive semidefinite. Hence, by Weyl's inequality and the triangle inequality,

$$\frac{2}{\pi} g(1) \|\Sigma\|_2 \leq \|T\|_2 \leq \frac{2}{\pi} \sum_{k=0}^{\infty} g(k) \|\Sigma \circ_k \Sigma\|_2. \tag{4.29}$$

The first half of inequality (4.29) yields the first half of inequality (2.8). Next, note that the $\Sigma \circ_k \Sigma$'s, in addition to being positive semidefinite, all have diagonal elements identically equal to one. Then, by [19], Theorem 5.5.18, we have for all $k \geq 2$, $\|\Sigma \circ_k \Sigma\|_2 = \|(\Sigma \circ_{k-1} \Sigma) \circ \Sigma\|_2 \leq \|\Sigma\|_2$. Therefore, the second half of inequality (4.29) yields

$$\|T\|_2 \leq \frac{2}{\pi} \sum_{k=0}^{\infty} g(k) \|\Sigma\|_2 = \|\Sigma\|_2,$$

which is the second half of inequality (2.8). □

4.4. Proof of Theorem 2.4

Because Σ belongs to the feasible region \mathcal{F} , and $\widehat{\Sigma}^{\text{generic}+}$ minimizes $\|\Sigma' - \widehat{\Sigma}^{\text{generic}}\|$ over $\Sigma' \in \mathcal{F}$ by (2.9), we conclude that

$$\|\widehat{\Sigma}^{\text{generic}+} - \widehat{\Sigma}^{\text{generic}}\| \leq \|\Sigma - \widehat{\Sigma}^{\text{generic}}\|. \tag{4.30}$$

Then, plugging inequality (4.30) into the triangle inequality

$$\|\widehat{\Sigma}^{\text{generic}+} - \Sigma\| \leq \|\widehat{\Sigma}^{\text{generic}+} - \widehat{\Sigma}^{\text{generic}}\| + \|\widehat{\Sigma}^{\text{generic}} - \Sigma\|$$

yields the conclusion of the theorem.

4.5. Proof of Corollary 2.5

First, with the choice $\mathcal{F} = \mathcal{S}_+^d$, we clearly have $\Sigma \in \mathcal{F}$. Then inequality (2.10) follows straightforwardly from Theorem 2.4. Next, we consider the choice of \mathcal{F} as in (2.11). With argument similar to that used in the proof of Lemma 4.4, we conclude that there exists an event A with probability at least $1 - \frac{1}{4}\alpha^2$ such that \widehat{T} satisfies

$$\|\widehat{T} - T\|_\infty \leq \sqrt{\frac{3}{2}}d^{-1/2}f(n, d, \alpha) \quad (4.31)$$

on the event A . For the rest of the proof, we concentrate on the event A . By (1.4), (1.5), (4.31) and the Lipschitz property of the sine function, we have

$$\|\widehat{\Sigma} - \Sigma\|_\infty \leq \frac{\pi}{2}\sqrt{\frac{3}{2}}d^{-1/2}f(n, d, \alpha) = C_3d^{-1/2}f(n, d, \alpha), \quad (4.32)$$

which further implies that $\Sigma \in \mathcal{F}$. Then inequality (2.10) again follows from Theorem 2.4. Finally, inequality (2.12) follows because $\|\widehat{\Sigma}^+ - \widehat{\Sigma}\|_\infty \leq C_3d^{-1/2}f(n, d, \alpha)$ by the choice (2.11) of \mathcal{F} , inequality (4.32), and the triangle inequality $\|\widehat{\Sigma}^+ - \Sigma\|_\infty \leq \|\widehat{\Sigma}^+ - \widehat{\Sigma}\|_\infty + \|\widehat{\Sigma} - \Sigma\|_\infty$.

5. Proof of Theorem 3.1

We first establish a proposition, which serves as the main ingredient for the proof of Theorem 3.1. For brevity of presentation, we denote

$$E = \widehat{\Sigma} - \Sigma.$$

Proposition 5.1. *Assume that Θ^* satisfies $0 < r < d$ and $\lambda_r(\Theta^*) \geq 2\mu$. On the event $\{2\|E\|_2 < \mu\}$, we have*

$$\widehat{r} = r, \quad (5.1)$$

$$\|\widehat{\Theta} - \Theta^*\|_F^2 \leq 8r\|E\|_2^2, \quad (5.2)$$

$$|\widehat{\sigma}^2 - \sigma^2| \leq \|E\|_2. \quad (5.3)$$

Proof. Let $\lambda_1(M) \geq \dots \geq \lambda_d(M)$ be the ordered eigenvalues of a generic symmetric matrix $M \in \mathbb{R}^{d \times d}$. Note that

$$\widehat{r} > r \iff \widehat{\lambda}_{r+1} - \widehat{\lambda}_d \geq \mu, \quad (5.4)$$

$$\widehat{r} < r \iff \widehat{\lambda}_r - \widehat{\lambda}_d < \mu. \quad (5.5)$$

We obtain, using Weyl's inequality,

$$\widehat{\lambda}_{r+1} - \widehat{\lambda}_d = \lambda_{r+1}(\Sigma + E) - \lambda_d(\Sigma + E) \leq \lambda_{r+1}(\Sigma) + 2\|E\|_2 - \lambda_d(\Sigma) = 2\|E\|_2, \quad (5.6)$$

$$\begin{aligned} \widehat{\lambda}_r - \widehat{\lambda}_d &= \lambda_r(\Sigma + E) - \lambda_d(\Sigma + E) \geq \lambda_r(\Sigma) - 2\|E\|_2 - \lambda_d(\Sigma) \\ &= \lambda_r(\Theta^*) - 2\|E\|_2. \end{aligned} \quad (5.7)$$

Together, (5.4), (5.5), (5.6), (5.7) and the condition $\lambda_r(\Theta^*) \geq 2\mu$ lead to

$$\{\widehat{r} \neq r\} \subseteq \{2\|E\|_2 \geq \min(\mu, \lambda_r(\Theta^*) - \mu)\} \subseteq \{2\|E\|_2 \geq \mu\}. \quad (5.8)$$

A similar reasoning is used in the proof of [3], Theorem 2. Consequently, equation (5.1), that is, $\widehat{r} = r$, holds on the event $\{2\|E\|_2 < \mu\}$, and for the rest of the proof we concentrate on this event. Then we have

$$\begin{aligned} \|\widehat{\Theta} - \Theta^*\|_F &\leq \sqrt{2r} \|\widehat{\Theta} - \Theta^*\|_2 = \sqrt{2r} \left\| \sum_{k=1}^r (\widehat{\lambda}_k - \widehat{\sigma}^2) \widehat{u}_k \widehat{u}_k^T - \Theta^* \right\|_2 \\ &= \sqrt{2r} \left\| \sum_{k=1}^d \widehat{\lambda}_k \widehat{u}_k \widehat{u}_k^T - \sum_{k=r+1}^d \widehat{\lambda}_k \widehat{u}_k \widehat{u}_k^T - \sum_{k=1}^r \widehat{\sigma}^2 \widehat{u}_k \widehat{u}_k^T - \Theta^* \right\|_2 \\ &= \sqrt{2r} \left\| \widehat{\Sigma} - \Sigma + \sigma^2 I_d - \sum_{k=1}^r \widehat{\sigma}^2 \widehat{u}_k \widehat{u}_k^T - \sum_{k=r+1}^d \widehat{\lambda}_k \widehat{u}_k \widehat{u}_k^T \right\|_2 \\ &= \sqrt{2r} \left\| E + \sum_{k=1}^d (\sigma^2 - \widetilde{\lambda}_k) \widehat{u}_k \widehat{u}_k^T \right\|_2 \\ &\leq \sqrt{2r} \left[\|E\|_2 + \max_{1 \leq k \leq d} |\widetilde{\lambda}_k - \sigma^2| \right]. \end{aligned} \quad (5.9)$$

Here, we have denoted

$$\widetilde{\lambda}_k = \begin{cases} \widehat{\sigma}^2, & \text{if } k \leq r, \\ \widehat{\lambda}_k, & \text{if } k \geq r+1. \end{cases}$$

We use Weyl's inequality again to observe that

$$\begin{aligned} \max_{1 \leq k \leq d} |\widetilde{\lambda}_k - \sigma^2| &= \max(|\lambda_{r+1}(\widehat{\Sigma}) - \sigma^2|, \dots, |\lambda_d(\widehat{\Sigma}) - \sigma^2|, |\widehat{\sigma}^2 - \sigma^2|) \\ &= \max(|\lambda_{r+1}(\widehat{\Sigma}) - \lambda_{r+1}(\Sigma)|, \dots, |\lambda_d(\widehat{\Sigma}) - \lambda_d(\Sigma)|) \\ &\leq \|E\|_2, \end{aligned} \quad (5.10)$$

which implies inequality (5.3). Finally, inequalities (5.9) and (5.10) together imply inequality (5.2). \square

Note that the regularization parameter μ should both be large enough such that the event $\{2\|E\|_2 < \mu\}$ has high probability, and be small enough such that the condition $\lambda_r(\Theta^*) \geq 2\mu$ is not too stringent. However, these requirements cannot always be met at the same time, as we demonstrate next. For brevity, we set $f = f(n, d, \alpha)$.

First, on the one hand, it is clear from Theorem 2.2 that we should choose, for some absolute constants c_1, c_2 and $\alpha < 1/2$,

$$\mu \approx c_1 \sqrt{\|T\|_2} f + c_2 f^2, \quad (5.11)$$

to guarantee that the event $\{2\|E\|_2 < \mu\}$ has probability larger than $1 - 2\alpha$. (In practice, we need a procedure that determines μ based on $\|\widehat{T}\|_2$ instead of $\|T\|_2$, and at the same time guarantees the convergence rates in (5.2) and (5.3) in terms of $\|T\|_2$. Theorem 3.1 describes such a procedure in detail, using the results from Theorem 2.2.) On the other hand, by Theorem 2.3 and the condition $\lambda_r(\Theta^*) \geq 2\mu$, the following string of inequalities

$$\frac{\pi}{2} \|T\|_2 \geq \|\Sigma\|_2 \geq \lambda_{\max}(\Theta^*) \geq \lambda_r(\Theta^*) \geq 2\mu \tag{5.12}$$

hold. Now, if $\|T\|_2 \ll f^2$, then $\mu \ll f^2$ as well by (5.12), contradicting (5.11). Therefore, the interesting case is (roughly) when inequality (2.2) holds.

Proof of Theorem 3.1. Let

$$\bar{\mu}' = 2\{C_1 \max[\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)] + (C_1 + C_2) f^2(n, d, \alpha)\}. \tag{5.13}$$

Then Theorem 2.2 guarantees that $\mathbb{P}\{2\|E\|_2 < \mu < \bar{\mu}'\} \geq 1 - \alpha - \alpha^2/4 > 1 - 2\alpha$ with the choices (3.4) and (5.13) of μ and $\bar{\mu}'$, and for the rest of the proof we concentrate on this event. Assume that Θ^* satisfies $0 < r < d$ and $\lambda_r(\Theta^*) \geq 2\bar{\mu}$, and n is large enough such that condition (2.2), which is in place for the reasons discussed in the remarks following Proposition 5.1, holds. Because condition (2.2) also ensures that equation (2.3) holds, we have $\bar{\mu}' = \bar{\mu}$. Hence, the assumption $\lambda_r(\Theta^*) \geq 2\bar{\mu}$ further implies that $\lambda_r(\Theta^*) \geq 2\bar{\mu}' > 2\mu$. Then Proposition 5.1 states that equation (3.6) and inequalities (5.2), (5.3) hold. Next, we can replace $\|E\|_2$ in inequalities (5.2) and (5.3) by $\bar{\mu}'/2$ using the bound $\|E\|_2 < \bar{\mu}'/2$, and further replace $\bar{\mu}'$ by $\bar{\mu}$. Inequality (3.8) and the second half of inequality (3.7) then follow. The first half of inequality (3.7) follows because by (3.3), we have

$$\|\tilde{\Sigma}^e - \Sigma\|_F^2 = \|\widehat{\Theta}_o - \Theta_o^*\|_F^2 \leq \|\widehat{\Theta} - \Theta^*\|_F^2.$$

It remains to establish the last statement of the theorem. We let $\text{diag}(\Theta^*)$ be the common value of the diagonal elements of Θ^* . We assume that $\text{diag}(\Theta^*) \leq 1 - \sqrt{2r\bar{\mu}^2}$ as in the statement of the theorem, and show that $\tilde{\Sigma}^e$ is positive semidefinite. Inequality (3.7) implies that $\|\widehat{\Theta} - \Theta^*\|_\infty \leq \sqrt{2r\bar{\mu}^2}$. Thus, the values of the diagonal elements of $\widehat{\Theta}$ cannot exceed $\text{diag}(\Theta^*) + \sqrt{2r\bar{\mu}^2} \leq 1$. Hence, in this case, by (3.3), $\tilde{\Sigma}^e$ is obtained by adding to $\widehat{\Theta}$ a diagonal matrix with nonnegative diagonal entries. Because $\widehat{\Theta}$ is positive semidefinite by construction, we conclude that $\tilde{\Sigma}^e$ is positive semidefinite as well. \square

6. Proof of Theorem 3.2

6.1. Preliminaries

We let $M \in \mathbb{R}^{d \times d}$ be an arbitrary matrix of rank r , with the (reduced) singular value decomposition $M = U \Lambda V^T$. Here, $U, V \in \mathbb{R}^{d \times r}$ are, respectively, matrix of the left and right orthonormal

singular vectors of M corresponding to the nonzero singular values that are the diagonal elements of $\Lambda \in \mathbb{R}^{r \times r}$. Following the exposition in [10], the tangent space $T(M) \subset \mathbb{R}^{d \times d}$ at M with respect to the algebraic variety of matrices with rank at most $r = \text{rank}(M)$, or the tangent space $T(M)$ for short, is given by

$$T(M) = \{UX^T + YV^T \mid X, Y \in \mathbb{R}^{d \times r}\}.$$

We denote the orthogonal complement of $T(M)$ by $T(M)^\perp$. In addition, we denote the projector onto the tangent space $T(M)$ by $\mathcal{P}_{T(M)}$, and the projector onto $T(M)^\perp$ by $\mathcal{P}_{T(M)^\perp}$. Then, for an arbitrary matrix $N \in \mathbb{R}^{d \times d}$, the explicit forms of $\mathcal{P}_{T(M)}$ and $\mathcal{P}_{T(M)^\perp}$ are given by

$$\begin{aligned} \mathcal{P}_{T(M)}(N) &= UU^T N + NVV^T - UU^T NVV^T, \\ \mathcal{P}_{T(M)^\perp}(N) &= (I_d - UU^T)N(I_d - VV^T), \end{aligned}$$

respectively. One basic fact involving the projectors $\mathcal{P}_{T(M)}$ and $\mathcal{P}_{T(M)^\perp}$ is

$$\|\mathcal{P}_{T(M)}(N)\|_2 \leq 2\|N\|_2 \quad \text{and} \quad \|\mathcal{P}_{T(M)^\perp}(N)\|_2 \leq \|N\|_2.$$

We denote the set of $d \times d$ diagonal matrices by Ω . We let the projector onto Ω be denoted by \mathcal{P}_Ω . Recall that \circ denotes the Hadamard product. Then, for an arbitrary matrix $N \in \mathbb{R}^{d \times d}$, the explicit form of \mathcal{P}_Ω is given by

$$\mathcal{P}_\Omega(N) = I_d \circ N.$$

We also prove a simple lemma.

Lemma 6.1. *Let $A, B, C \in \mathbb{R}^{d \times d}$ be arbitrary matrices. Then*

$$\|ACB\|_\infty \leq \sqrt{\|AA^T\|_\infty \|B^T B\|_\infty} \|C\|_2.$$

Proof. The proof can be found in Appendix B. □

6.2. Recovery bound with primal-dual certificate

We let $\bar{\Theta}, Q \in \mathbb{R}^{d \times d}$ but otherwise be arbitrary at this stage. Eventually, we will set $\bar{\Theta}$ to be some low-rank approximation to Θ^* , and set Q to be a *primal-dual certificate* [48], or certificate for short, in the sense defined in equation (6.12) below. For notational brevity, we denote

$$\bar{T} = T(\bar{\Theta}) \quad \text{and} \quad \bar{T}^\perp = T(\bar{\Theta})^\perp$$

for the tangent space $T(\bar{\Theta})$ and its orthogonal complement $T(\bar{\Theta})^\perp$, respectively.

We now state two lemmas toward the general recovery bound for the refined estimator $\tilde{\Sigma}$ in terms of $\bar{\Theta}$ and the (soon-to-be) certificate Q .

Lemma 6.2. *We have*

$$\begin{aligned} & \frac{1}{2} \|\tilde{\Theta}_o - \Theta_o^*\|_F^2 + \frac{1}{2} \|\tilde{\Theta}_o - Q_o\|_F^2 + \langle -Q_o + \bar{\Theta}_o - \Theta_o^* + \tilde{\Theta}_o, \bar{\Theta}_o - \tilde{\Theta}_o \rangle \\ & = \frac{1}{2} \|\bar{\Theta}_o - \Theta_o^*\|_F^2 + \frac{1}{2} \|\bar{\Theta}_o - Q_o\|_F^2. \end{aligned} \quad (6.1)$$

Proof. The identity follows from straightforward algebra, and can also be obtained from the proof for [48], Theorem 3.2. \square

We define, for any constant $c \geq 1$,

$$G_c = \{ \Phi \in \mathbb{R}^{d \times d}: \Phi \in \mu \partial \|\bar{\Theta}\|_* \text{ and } \|\mathcal{P}_{\bar{T}^\perp} \Phi\|_2 \leq \mu/c \}. \quad (6.2)$$

Here, $\partial \|A\|_*$ denotes the subdifferential with respect to the nuclear norm at the matrix A ; we refer to [44] for its explicit form. Note that G_c is a subset of the subdifferential $\mu \partial \|\bar{\Theta}\|_*$, and coincides with the latter when $c = 1$.

Lemma 6.3. *Assume that*

$$-Q_o + \bar{\Theta}_o - \Theta_o^* + \widehat{\Sigma}_o \in G_c. \quad (6.3)$$

Then

$$\langle -Q_o + \bar{\Theta}_o - \Theta_o^* + \tilde{\Theta}_o, \bar{\Theta}_o - \tilde{\Theta}_o \rangle \geq (1 - 1/c) \mu \|\mathcal{P}_{\bar{T}^\perp} \tilde{\Theta}\|_*. \quad (6.4)$$

Proof. We follow the proof of [48], Proposition 3.2. Let $\Psi, \Xi \in \mathbb{R}^{d \times d}$ satisfy $\Psi \in \mu \partial \|\tilde{\Theta}\|_*$, $\Xi_o \in \mu \partial \|\bar{\Theta}\|_*$ but otherwise be arbitrary at this stage. By the definition of subgradient, we have

$$\langle \Xi_o, \bar{\Theta} - \tilde{\Theta} \rangle \geq \mu \|\bar{\Theta}\|_* - \mu \|\tilde{\Theta}\|_* \geq \langle \Psi, \bar{\Theta} - \tilde{\Theta} \rangle. \quad (6.5)$$

Now we impose on Ξ the stronger condition that $\Xi_o \in G_c$. Then the first half of inequality (6.5) can be strengthened by [20], Lemma 6, to

$$\langle \Xi_o, \bar{\Theta} - \tilde{\Theta} \rangle \geq (1 - 1/c) \mu \|\mathcal{P}_{\bar{T}^\perp} \tilde{\Theta}\|_* + \mu \|\bar{\Theta}\|_* - \mu \|\tilde{\Theta}\|_*. \quad (6.6)$$

Next, combining inequality (6.6) and the second half of inequality (6.5) yields

$$\langle \Xi_o, \bar{\Theta} - \tilde{\Theta} \rangle \geq \langle \Psi, \bar{\Theta} - \tilde{\Theta} \rangle + (1 - 1/c) \mu \|\mathcal{P}_{\bar{T}^\perp} \tilde{\Theta}\|_*. \quad (6.7)$$

Let $L(\Theta) = \frac{1}{2} \|\Theta_o - \widehat{\Sigma}_o\|_F^2$ denote the loss function in the convex program (3.13) and $\nabla L(\Theta) = \Theta_o - \widehat{\Sigma}_o$ denote its gradient. Then, adding $\langle \nabla L(\tilde{\Theta}), \bar{\Theta} - \tilde{\Theta} \rangle$ to both sides of inequality (6.7) yields

$$\langle \Xi_o + \nabla L(\tilde{\Theta}), \bar{\Theta} - \tilde{\Theta} \rangle \geq \langle \Psi + \nabla L(\tilde{\Theta}), \bar{\Theta} - \tilde{\Theta} \rangle + (1 - 1/c) \mu \|\mathcal{P}_{\bar{T}^\perp} \tilde{\Theta}\|_*. \quad (6.8)$$

We now fix our choices of Ψ and Ξ . First, by the optimality of $\tilde{\Theta}$ for the convex program (3.13), we have $0 \in \nabla L(\tilde{\Theta}) + \mu \partial \|\tilde{\Theta}\|_*$. Hence, we can fix $\Psi \in \mu \partial \|\tilde{\Theta}\|_*$ such that

$$\nabla L(\tilde{\Theta}) + \Psi = 0. \quad (6.9)$$

Then, plugging equation (6.9) into inequality (6.8) yields

$$\langle \Xi_o + \nabla L(\tilde{\Theta}), \bar{\Theta} - \tilde{\Theta} \rangle \geq (1 - 1/c)\mu \|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Theta}\|_*. \tag{6.10}$$

Next, we set $\Xi = -Q + \bar{\Theta} - \Theta^* + \hat{\Sigma}$, so $\Xi_o \in G_c$ by assumption. We also use $\nabla L(\tilde{\Theta}) = \tilde{\Theta}_o - \hat{\Sigma}_o$. Then inequality (6.10) becomes

$$\langle -Q_o + \bar{\Theta}_o - \Theta_o^* + \tilde{\Theta}_o, \bar{\Theta} - \tilde{\Theta} \rangle \geq (1 - 1/c)\mu \|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Theta}\|_*. \tag{6.11}$$

Finally, observe that, for arbitrary commensurate matrices A and B , we have $\langle A_o, B \rangle = \text{tr}(A_o^T B) = \text{tr}(A_o^T B_o) = \langle A_o, B_o \rangle$. Hence, we are free to replace the term $\bar{\Theta} - \tilde{\Theta}$ in the angle bracket on the left-hand side of inequality (6.11) by $\bar{\Theta}_o - \tilde{\Theta}_o$. The corollary then follows. \square

We are now ready to derive the general recovery bound for the refined estimator $\tilde{\Sigma}$ in terms of $\bar{\Theta}$ and the certificate Q . We denote $E = \hat{\Sigma} - \Sigma$ again, and note that $E_o = E$.

Theorem 6.4. *If*

$$-Q_o + \bar{\Theta}_o + E \in G_c, \tag{6.12}$$

then

$$\frac{1}{2} \|\tilde{\Sigma} - \Sigma\|_F^2 + (1 - 1/c)\mu \|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Theta}\|_* \leq \frac{1}{2} \|\bar{\Theta}_o - \Theta_o^*\|_F^2 + \frac{1}{2} \|\bar{\Theta}_o - Q_o\|_F^2. \tag{6.13}$$

Proof. We start from Lemma 6.2. By the construction of $\tilde{\Sigma}$ as in (3.12), the off-diagonal elements of $\tilde{\Theta}$ and $\tilde{\Sigma}$ agree, that is, $\tilde{\Theta}_o = \tilde{\Sigma}_o$. In addition, $\Theta_o^* = \Sigma_o$. Hence, $\bar{\Theta}_o - \Theta_o^* = \tilde{\Sigma}_o - \Sigma_o = \tilde{\Sigma} - \Sigma$. Thus, after discarding the term $\frac{1}{2} \|\tilde{\Theta}_o - Q_o\|_F^2$, equation (6.1) becomes

$$\begin{aligned} & \frac{1}{2} \|\tilde{\Sigma} - \Sigma\|_F^2 + \langle -Q_o + \bar{\Theta}_o - \Theta_o^* + \tilde{\Theta}_o, \bar{\Theta}_o - \tilde{\Theta}_o \rangle \\ & \leq \frac{1}{2} \|\bar{\Theta}_o - \Theta_o^*\|_F^2 + \frac{1}{2} \|\bar{\Theta}_o - Q_o\|_F^2. \end{aligned} \tag{6.14}$$

Next we invoke Lemma 6.3. Because $-\Theta_o^* + \hat{\Sigma}_o = -\Sigma_o + \hat{\Sigma}_o = E_o = E$, condition (6.12) translates into condition (6.3), and hence inequality (6.4) holds. Finally, plugging inequality (6.4) into inequality (6.14) yields the theorem. \square

6.3. Certificate construction

From Theorem 6.4, it is clear that the recovery bounds on $\|\tilde{\Sigma} - \Sigma\|_F^2$ and $\|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Theta}\|_*$ depend crucially on an appropriate certificate Q such that $\|Q_o - \bar{\Theta}_o\|_F^2$ can be tightly bounded. This section is dedicated to the construction of such a certificate.

Recall that $\bar{\Theta} \in \mathbb{R}^{d \times d}$, which is intended to be some low-rank approximation to Θ^* , has been left unspecified so far. Now we restrict $\bar{\Theta}$ to be a positive semidefinite matrix of rank r , with the eigen-decomposition

$$\bar{\Theta} = \bar{U} \bar{\Lambda} \bar{U}^T. \tag{6.15}$$

Here, $\bar{U} \in \mathbb{R}^{d \times r}$ is the matrix of the orthonormal eigenvectors of $\bar{\Theta}$ corresponding to the positive eigenvalues that are the diagonal elements of $\bar{\Lambda} \in \mathbb{R}^{r \times r}$. Recall from Section 6.2 that \bar{T} denotes the tangent space $T(\bar{\Theta})$, and \bar{T}^\perp denotes its orthogonal complement $T(\bar{\Theta})^\perp$. Then, with our specific choice of $\bar{\Theta}$, the projectors $\mathcal{P}_{\bar{T}}$ and $\mathcal{P}_{\bar{T}^\perp}$ are given by

$$\mathcal{P}_{\bar{T}}(N) = \bar{U}\bar{U}^T N + N\bar{U}\bar{U}^T - \bar{U}\bar{U}^T N\bar{U}\bar{U}^T, \quad (6.16a)$$

$$\mathcal{P}_{\bar{T}^\perp}(N) = (I_d - \bar{U}\bar{U}^T)N(I_d - \bar{U}\bar{U}^T) \quad (6.16b)$$

for arbitrary $N \in \mathbb{R}^{d \times d}$. For notational brevity, from now on we will omit the parentheses surrounding the argument when applying the projectors. Again with our specific choice of $\bar{\Theta}$, we can give a more explicit characterization of G_c , defined earlier in (6.2), as

$$G_c = \{\Phi \in \mathbb{R}^{d \times d}: \mathcal{P}_{\bar{T}}\Phi = \mu\bar{U}\bar{U}^T \text{ and } \|\mathcal{P}_{\bar{T}^\perp}\Phi\|_2 \leq \mu/c\}. \quad (6.17)$$

We also define

$$\gamma = \|\bar{U}\bar{U}^T\|_\infty = \max_{1 \leq i \leq d} [\bar{U}\bar{U}^T]_{ii} \leq 1. \quad (6.18)$$

The second equality in (6.18) is due to the fact that $\bar{U}\bar{U}^T$ is positive semidefinite, while the inequality follows since \bar{U} is a matrix of orthonormal eigenvectors.

Next, we obtain some technical results stating that, under certain conditions, the operators $\mathcal{P}_{\bar{T}}$ and $\mathcal{P}_\Omega \mathcal{P}_{\bar{T}}$ are contractions under certain matrix norms (Lemma 6.5), and the operator $\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega$, with \mathcal{I}_d the identity operator in $\mathbb{R}^{d \times d}$, is invertible (Lemma 6.6). These results essentially follow from [20] (e.g., their Lemmas 4, 8 and 10), but we offer tighter bounds specialized to our study.

Lemma 6.5. *For any diagonal matrix $D \in \mathbb{R}^{d \times d}$, we have*

$$\|\mathcal{P}_{\bar{T}}D\|_\infty \leq 3\gamma\|D\|_\infty. \quad (6.19)$$

For any matrix $M \in \mathbb{R}^{d \times d}$, we have

$$\|\mathcal{P}_{\bar{T}}M\|_\infty \leq 2\sqrt{\gamma}\|M\|_2 \quad (6.20)$$

and

$$\|\mathcal{P}_\Omega \mathcal{P}_{\bar{T}}M\|_1 \leq 3\gamma\|M\|_1. \quad (6.21)$$

Proof. The proof can be found in Appendix B. □

Lemma 6.6. *Assume that $\gamma < 1/3$. Then the operator $\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega: \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d \times d}$ is a bijection, and hence is invertible. Moreover, $\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega$ satisfies, for any matrix $M \in \mathbb{R}^{d \times d}$,*

$$\|(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)^{-1}M\|_\infty \leq \frac{1}{1 - 3\gamma}\|M\|_\infty. \quad (6.22)$$

Proof. The proof can be found in Appendix B. □

We demonstrate in Theorem 6.7 that, under appropriate conditions, we can solve for $Q_o - \bar{\Theta}_o$ in an equation of the form (6.12), such that $Q - \bar{\Theta}$ has low rank and $\|Q - \bar{\Theta}\|_2$ is small, which further implies that $\|Q_o - \bar{\Theta}_o\|_F^2$ is tightly bounded, as is desired. The techniques we use are based on the proofs of [9], Proposition 5.2 and [20], Theorem 5.

Theorem 6.7. *Assume that $\bar{\Theta}$ is positive semidefinite and has the eigen-decomposition (6.15). Let $\bar{T} = T(\bar{\Theta})$. Let G_c and γ be defined as in (6.17) and (6.18), respectively. Suppose that γ satisfies*

$$\gamma < \frac{1}{c + 3}. \tag{6.23}$$

Let A be the event on which

$$\mu \geq \left(\frac{1}{c} - \frac{\gamma}{1 - 3\gamma}\right)^{-1} \left(\frac{2\sqrt{\gamma}}{1 - 3\gamma} + 1\right) \|E\|_2 \tag{6.24}$$

holds. Then, on the event A , there exists some $\Phi \in \bar{T}$ such that

$$-\Phi_o + E \in G_c \tag{6.25}$$

and

$$\|\Phi\|_2 \leq \left(\frac{2}{c} + 1\right) \mu. \tag{6.26}$$

Remark. Note that inequality (6.23) ensures that the multiplicative factor $\left(\frac{1}{c} - \frac{\gamma}{1 - 3\gamma}\right)^{-1}$ in inequality (6.24) is positive.

Proof of Theorem 6.7. We focus on the event A . Note that assumption (6.23) entails that $\gamma < 1/4$ since $c \geq 1$. As a result, we can apply Lemma 6.6 to conclude that $\mathcal{I}_d - \mathcal{P}_{\bar{T}}\mathcal{P}_{\Omega}$ is invertible, and that inequality (6.22) holds. Then we can set

$$\Phi = (\mathcal{I}_d - \mathcal{P}_{\bar{T}}\mathcal{P}_{\Omega})^{-1} (\mathcal{P}_{\bar{T}}E - \mu\bar{U}\bar{U}^T). \tag{6.27}$$

We show that Φ has all the desired properties.

First, we apply the operator $\mathcal{I}_d - \mathcal{P}_{\bar{T}}\mathcal{P}_{\Omega}$ on both sides of equation (6.27), and obtain

$$\Phi = \mathcal{P}_{\bar{T}}\mathcal{P}_{\Omega}\Phi + \mathcal{P}_{\bar{T}}E - \mu\bar{U}\bar{U}^T, \tag{6.28}$$

from which it is clear that $\Phi \in \bar{T}$.

Relationship (6.25) is equivalent to

$$-(\Phi - \mathcal{P}_{\Omega}\Phi) + E \in G_c, \tag{6.29}$$

which is further equivalent to the following two conditions by the characterization (6.17) of G_c . The first condition is obtained by applying the operator $\mathcal{P}_{\bar{\tau}}$ and the second one is obtained by applying the operator $\mathcal{P}_{\bar{\tau}^\perp}$ on both sides of (6.29):

$$-(\mathcal{I}_d - \mathcal{P}_{\bar{\tau}}\mathcal{P}_\Omega)\Phi + \mathcal{P}_{\bar{\tau}}E = \mu\bar{U}\bar{U}^T, \quad (6.30a)$$

$$\|\mathcal{P}_{\bar{\tau}^\perp}(\Phi - \mathcal{P}_\Omega\Phi - E)\|_2 \leq \mu/c. \quad (6.30b)$$

Equation (6.30a) is equivalent to equation (6.28), and hence is satisfied. Next, we check that inequality (6.30b) holds. By equation (6.27), inequalities (6.22) and (6.20), we have

$$\begin{aligned} \|\Phi\|_\infty &\leq \frac{1}{1-3\gamma} \|\mathcal{P}_{\bar{\tau}}E - \mu\bar{U}\bar{U}^T\|_\infty \leq \frac{1}{1-3\gamma} (\|\mathcal{P}_{\bar{\tau}}E\|_\infty + \|\mu\bar{U}\bar{U}^T\|_\infty) \\ &\leq \frac{1}{1-3\gamma} (2\sqrt{\gamma}\|E\|_2 + \gamma\mu). \end{aligned} \quad (6.31)$$

Using inequality (6.31) and $\|\mathcal{P}_{\bar{\tau}^\perp}\mathcal{P}_\Omega\Phi\|_2 \leq \|\mathcal{P}_\Omega\Phi\|_2 = \|\mathcal{P}_\Omega\Phi\|_\infty \leq \|\Phi\|_\infty$, we have

$$\begin{aligned} \|\mathcal{P}_{\bar{\tau}^\perp}(\Phi - \mathcal{P}_\Omega\Phi - E)\|_2 &\leq \|\mathcal{P}_{\bar{\tau}^\perp}\Phi\|_2 + \|\mathcal{P}_{\bar{\tau}^\perp}\mathcal{P}_\Omega\Phi\|_2 + \|\mathcal{P}_{\bar{\tau}^\perp}E\|_2 \\ &\leq 0 + \|\Phi\|_\infty + \|E\|_2 \leq \left(\frac{2\sqrt{\gamma}}{1-3\gamma} + 1\right)\|E\|_2 + \frac{\gamma}{1-3\gamma}\mu. \end{aligned} \quad (6.32)$$

Then it is easy to see that inequality (6.32), assumptions (6.23) and (6.24) together imply inequality (6.30b). Hence, we have verified (6.25).

Finally, starting from equation (6.28), we have

$$\begin{aligned} \|\Phi\|_2 &\leq \|\mathcal{P}_{\bar{\tau}}\mathcal{P}_\Omega\Phi\|_2 + \|\mathcal{P}_{\bar{\tau}}E\|_2 + \|\mu\bar{U}\bar{U}^T\|_2 \leq 2\|\Phi\|_\infty + 2\|E\|_2 + \mu\|\bar{U}\bar{U}^T\|_2 \\ &\leq \frac{2}{1-3\gamma} (2\sqrt{\gamma}\|E\|_2 + \gamma\mu) + 2\|E\|_2 + \mu = 2\left(\frac{2\sqrt{\gamma}}{1-3\gamma} + 1\right)\|E\|_2 + \left(\frac{2\gamma}{1-3\gamma} + 1\right)\mu \\ &\leq 2\left(\frac{1}{c} - \frac{\gamma}{1-3\gamma}\right)\mu + \left(\frac{2\gamma}{1-3\gamma} + 1\right)\mu = \left(\frac{2}{c} + 1\right)\mu. \end{aligned}$$

Here, the second inequality follows from the fact that $\|\mathcal{P}_{\bar{\tau}}\mathcal{P}_\Omega\Phi\|_2 \leq 2\|\mathcal{P}_\Omega\Phi\|_2 \leq 2\|\Phi\|_\infty$, the third inequality follows from inequality (6.31), and the fourth inequality follows by assumption (6.24). Hence, inequality (6.26) is established. \square

6.4. Recovery bound for the refined estimator $\tilde{\Sigma}$

In this section, we state in Corollary 6.8 the main recovery bound that will lead to the oracle inequality for the refined estimator $\tilde{\Sigma}$. We recall U_r^* , γ_r and Θ_r^* as introduced in equations (3.9), (3.10) and (3.11).

Corollary 6.8. *Let r be such that $0 \leq r \leq r^*$ and*

$$\gamma_r < \frac{1}{c+3}. \quad (6.33)$$

Let A be the event on which the regularization parameter μ satisfies

$$\mu \geq \left(\frac{1}{c} - \frac{\gamma_r}{1-3\gamma_r} \right)^{-1} \left(\frac{2\sqrt{\gamma_r}}{1-3\gamma_r} + 1 \right) \|E\|_2. \quad (6.34)$$

(Note that inequalities (6.33) and (6.34) are just inequalities (6.23) and (6.24) with the substitution of γ by γ_r .) Then, on the event A we have

$$\|\tilde{\Sigma} - \Sigma\|_F^2 + (2 - 2/c)\mu \|\mathcal{P}_{T(\Theta^*)^\perp} \tilde{\Theta}\|_* \leq \sum_{j: r < j \leq r^*} \lambda_j^2(\Theta^*) + 2(1 + 2/c)^2 r \mu^2. \quad (6.35)$$

Remark. We can now see that the choice $c = 1$ in G_c is sufficient for proving a bound on $\|\tilde{\Sigma} - \Sigma\|_F^2$. With this choice of c , inequality (6.33) states that U_r^* , the truncated matrix of the orthonormal eigenvectors of Θ^* corresponding to the r largest eigenvalues, should satisfy the mild condition $\|U_r^* U_r^{*T}\|_\infty < 1/4$. On the other hand, the choice $c > 1$ leads to a bound on $\|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1$ as we will see in Appendix C.

Proof of Corollary 6.8. We start with the general recovery bound, Theorem 6.4. In the context of Theorem 6.4, $\tilde{\Theta}$ and Q should satisfy relationship (6.12) but are otherwise completely arbitrary.

We now set $\tilde{\Theta} = \Theta_r^*$, so $\tilde{\Theta}$ is positive semidefinite. We also concentrate on the event A . Then, by assumptions (6.33) and (6.34), inequalities (6.23) and (6.24) hold with the substitution of γ by γ_r . Hence, Theorem 6.7 applies. We let Φ be constructed according to Theorem 6.7 for the chosen $\tilde{\Theta} = \Theta_r^*$, so that $\Phi \in \tilde{T} = T(\Theta_r^*)$, $-\Phi_o + E \in G_c$, and $\|\Phi\|_2 \leq (1 + 2/c)\mu$. We set $Q = \tilde{\Theta} + \Phi$ so $Q - \tilde{\Theta} = \Phi$. Then relationship (6.12) is satisfied, and Theorem 6.4 further states that inequality (6.13) holds. We proceed to bound the two terms on the right-hand side of inequality (6.13) separately.

First, we consider the term $\|\tilde{\Theta}_o - \Theta_o^*\|_F^2$. Here and below, for brevity, we sometimes abbreviate the summation range $j: r < j \leq r^*$ by $j > r$. We have

$$\|\tilde{\Theta}_o - \Theta_o^*\|_F^2 \leq \|\tilde{\Theta} - \Theta^*\|_F^2 = \|\Theta_r^* - \Theta^*\|_F^2 = \sum_{j>r} \lambda_j^2(\Theta^*).$$

Next, we consider the term $\|\tilde{\Theta}_o - Q_o\|_F^2$. Using the fact that $\Phi \in T(\Theta_r^*)$ and so $\text{rank}(\Phi) \leq 2r$, and $\|\Phi\|_2 \leq (1 + 2/c)\mu$, we have

$$\|\tilde{\Theta}_o - Q_o\|_F^2 = \|\Phi_o\|_F^2 \leq \|\Phi\|_F^2 \leq 2r \|\Phi\|_2^2 \leq 2(1 + 2/c)^2 r \mu^2.$$

Combining both displays, we conclude that inequality (6.35) holds. \square

The bound on $\|\tilde{\Sigma} - \Sigma\|_F$ obtained in Corollary 6.8 can be further refined by optimizing the balance between the approximation error and the estimation error. We can also fix our choice of

the regularization parameter μ according to inequality (6.34). These considerations finally lead to our proof of Theorem 3.2.

Proof of Theorem 3.2. We fix $c = 2$, and $\gamma' = 1/9$. Then inequality (6.33) holds with the substitution of γ_r by γ' . Let A be the event

$$A = \left\{ \left(\frac{1}{c} - \frac{\gamma'}{1 - 3\gamma'} \right)^{-1} \left(\frac{2\sqrt{\gamma'}}{1 - 3\gamma'} + 1 \right) \|E\|_2 \leq \mu \leq \bar{\mu} \right\}. \quad (6.36)$$

That is, A is the event on which both $\mu \leq \bar{\mu}$ and inequality (6.34) with the substitution of γ_r by γ' hold. Note that the multiplicative factor in front of $\|E\|_2$ on the right-hand side of (6.36) exactly equals $C = 6$ with our choices of c and γ' . Then, by Theorem 2.2 and our choices (3.15) and (3.16) of μ and $\bar{\mu}$, we conclude that $\mathbb{P}(A) \geq 1 - \alpha - \alpha^2/4 > 1 - 2\alpha$, and for the rest of the proof we concentrate on the event A .

We let R be chosen according to (3.17), so in particular $\gamma_R \leq 1/9 = \gamma'$. Because γ_r is non-decreasing in r , and inequalities (6.33) and (6.34) hold with the substitution of γ_r by γ' , it is straightforward to conclude that inequalities (6.33) and (6.34) hold in terms of γ_r for all $0 \leq r \leq R$. Hence, by Corollary 6.8, inequality (6.35) holds for all $0 \leq r \leq R$. Then, after discarding the term $(2 - 2/c)\mu \|\mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta}\|_*$ on the left-hand side of inequality (6.35), we obtain, for all $0 \leq r \leq R$, that

$$\|\tilde{\Sigma} - \Sigma\|_F^2 \leq \sum_{j>r} \lambda_j^2(\Theta^*) + 2(1 + 2/c)^2 r \mu^2 \leq \sum_{j>r} \lambda_j^2(\Theta^*) + 8r \bar{\mu}^2. \quad (6.37)$$

Here, the second inequality in (6.37) follows because $c = 2$ and $\mu \leq \bar{\mu}$. Finally, the theorem follows by taking the minimum of inequality (6.37) over $0 \leq r \leq R$. \square

Appendix A: Discussion of some basic concepts

In this section, we present formal definitions of some basic concepts in this paper and then discuss the characterization of the semi-parametric elliptical copula model. We first present the definition of an elliptical distribution; see, for instance, [7].

Definition A.1. A random vector $Y = (Y_1, \dots, Y_d)^T \in \mathbb{R}^d$ has an elliptical distribution if for some $\mu \in \mathbb{R}^d$ and some positive semidefinite matrix $\bar{\Sigma} \in \mathbb{R}^{d \times d}$, the characteristic function $\varphi_{Y-\mu}(t)$ of $Y - \mu$ is a function of the quadratic form $t^T \bar{\Sigma} t$, that is, $\varphi_{Y-\mu}(t) = \phi(t^T \bar{\Sigma} t)$ for some function ϕ . We write $Y \sim \mathcal{E}_d(\mu, \bar{\Sigma}, \phi)$, and call ϕ the characteristic generator.

Next, we present the definition of a copula [40]; see, for instance, [13], Theorem 2.2.

Definition A.2. The copula $C : [0, 1]^d \rightarrow [0, 1]$ of a continuous random vector $Y = (Y_1, \dots, Y_d)^T \in \mathbb{R}^d$ is the joint distribution function of the transformed random vector $U = (F_1(Y_1), \dots, F_d(Y_d))^T \in \mathbb{R}^d$ on the unit cube $[0, 1]^d$, using the marginal distribution functions $F_j(y) = \mathbb{P}\{Y_j \leq y\}$ for $1 \leq j \leq d$.

We recall the basic property that copulas are invariant under strictly increasing transformations of the individual vector components of the underlying distribution; see, for instance, [13], Theorem 2.6. It follows from this invariance property that, if the random vector $X \in \mathbb{R}^d$ follows a distribution from the semi-parametric elliptical copula model, and if X has the same copula with an elliptically distributed random vector $Y \in \mathbb{R}^d$ such that $Y \sim \mathcal{E}_d(\mu, \bar{\Sigma}, \phi)$, then the copula of X is uniquely characterized by the same characteristic generator ϕ and a copula correlation matrix Σ , defined as $[\Sigma]_{k\ell} = [\bar{\Sigma}]_{k\ell} / ([\bar{\Sigma}]_{kk}[\bar{\Sigma}]_{\ell\ell})^{1/2}$ for all $1 \leq k, \ell \leq d$.

Appendix B: Auxiliary proofs for Section 6

This section contains the proofs of some auxiliary lemmas in Section 6.

Proof of Lemma 6.1. We let $e_i \in \mathbb{R}^d$ denote the vector with one at the i th position and zeros elsewhere, and $\|\cdot\|$ denote the Euclidean norm for vectors. Then we have

$$\begin{aligned} \|ACB\|_\infty &= \max_{i,j} |e_i^T ACBe_j| \leq \max_{i,j} \|e_i^T A\| \|CBe_j\| \leq \max_{i,j} \|e_i^T A\| \|C\|_2 \|Be_j\| \\ &= \max_{i,j} \sqrt{e_i^T AA^T e_i} \|C\|_2 \sqrt{e_j^T B^T Be_j} \leq \sqrt{\|AA^T\|_\infty} \sqrt{\|B^T B\|_\infty} \|C\|_2. \end{aligned}$$

Here, the first equality follows from an observation in the proof of [10], Proposition 4, and the first inequality follows by the Cauchy–Schwarz inequality. The lemma follows. \square

Proof of Lemma 6.5. Let $D \in \mathbb{R}^{d \times d}$ be an arbitrary diagonal matrix, and $M \in \mathbb{R}^{d \times d}$ an arbitrary matrix. We first prove inequality (6.19). Using equation (6.16a), we have

$$\|\mathcal{P}_{\bar{T}} D\|_\infty \leq \|(\bar{U}\bar{U}^T)D\|_\infty + \|D(\bar{U}\bar{U}^T)\|_\infty + \|(\bar{U}\bar{U}^T)D(\bar{U}\bar{U}^T)\|_\infty. \tag{B.1}$$

We bound the terms on the right-hand side of inequality (B.1) separately. Note that, although $\|\cdot\|_\infty$, the element-wise ℓ_∞ norm, is not submultiplicative, it is easy to see that the inequality $\|AB\|_\infty \leq \|A\|_\infty \|B\|_\infty$ holds when at least one of A, B is a diagonal matrix. Hence, we have

$$\max\{\|(\bar{U}\bar{U}^T)D\|_\infty, \|D(\bar{U}\bar{U}^T)\|_\infty\} \leq \|\bar{U}\bar{U}^T\|_\infty \|D\|_\infty = \gamma \|D\|_\infty. \tag{B.2}$$

Next, setting $A = B = \bar{U}\bar{U}^T$ and $C = D$ in Lemma 6.1 yields

$$\begin{aligned} \|(\bar{U}\bar{U}^T)D(\bar{U}\bar{U}^T)\|_\infty &\leq \sqrt{\|\bar{U}\bar{U}^T \bar{U}\bar{U}^T\|_\infty} \sqrt{\|\bar{U}\bar{U}^T \bar{U}\bar{U}^T\|_\infty} \|D\|_2 \\ &= \sqrt{\|\bar{U}\bar{U}^T\|_\infty} \sqrt{\|\bar{U}\bar{U}^T\|_\infty} \|D\|_2 = \gamma \|D\|_\infty. \end{aligned} \tag{B.3}$$

Here, the final equality follows because D is diagonal and so $\|D\|_2 = \|D\|_\infty$. Finally, plugging inequalities (B.2) and (B.3) into inequality (B.1) yields inequality (6.19).

To prove inequality (6.20), note that, again by equation (6.16a), we have

$$\|\mathcal{P}_{\bar{T}} M\|_\infty \leq \|(\bar{U}\bar{U}^T)M\|_\infty + \|(I_d - \bar{U}\bar{U}^T)M(\bar{U}\bar{U}^T)\|_\infty. \tag{B.4}$$

Setting $A = UU^T$, $B = I_d$ and $C = M$ in Lemma 6.1 yields

$$\|(\bar{U}\bar{U}^T)M\|_\infty \leq \sqrt{\gamma}\|M\|_2, \quad (\text{B.5})$$

while setting $A = (I_d - \bar{U}\bar{U}^T)$, $B = UU^T$ and $C = M$ in Lemma 6.1 yields

$$\|(I_d - \bar{U}\bar{U}^T)M(\bar{U}\bar{U}^T)\|_\infty \leq \sqrt{\gamma}\|M\|_2. \quad (\text{B.6})$$

Inequality (6.20) then follows from inequalities (B.4), (B.5) and (B.6).

Finally, we prove inequality (6.21). Note that $\|\cdot\|_\infty$ and $\|\cdot\|_1$ are dual norms. Then

$$\begin{aligned} \|\mathcal{P}_\Omega \mathcal{P}_{\bar{T}} M\|_1 &= \sup_{N: \|N\|_\infty \leq 1} \langle \mathcal{P}_\Omega \mathcal{P}_{\bar{T}} M, N \rangle = \sup_{N: \|N\|_\infty \leq 1} \langle \mathcal{P}_{\bar{T}} M, \mathcal{P}_\Omega N \rangle = \sup_{N: \|N\|_\infty \leq 1} \langle M, \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega N \rangle \\ &\leq \sup_{N: \|N\|_\infty \leq 1} \|M\|_1 \|\mathcal{P}_{\bar{T}} \mathcal{P}_\Omega N\|_\infty \leq 3\gamma \sup_{N: \|N\|_\infty \leq 1} \|M\|_1 \|\mathcal{P}_\Omega N\|_\infty \\ &\leq 3\gamma \sup_{N: \|N\|_\infty \leq 1} \|M\|_1 \|N\|_\infty \leq 3\gamma \|M\|_1, \end{aligned}$$

using first Hölder's inequality and then inequality (6.19) on the diagonal matrix $\mathcal{P}_\Omega N$. \square

Proof of Lemma 6.6. We assume that $\gamma < 1/3$. Let $M \in \mathbb{R}^{d \times d}$ be an arbitrary matrix. Applying inequality (6.19) in Lemma 6.5 on the diagonal matrix $\mathcal{P}_\Omega M$, we obtain

$$\|\mathcal{P}_{\bar{T}} \mathcal{P}_\Omega M\|_\infty \leq 3\gamma \|\mathcal{P}_\Omega M\|_\infty \leq 3\gamma \|M\|_\infty.$$

Then, by the triangle inequality,

$$\|(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)M\|_\infty \geq \|M\|_\infty - \|\mathcal{P}_{\bar{T}} \mathcal{P}_\Omega M\|_\infty \geq (1 - 3\gamma)\|M\|_\infty.$$

Because $\gamma < 1/3$, $\|(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)M\|_\infty = 0$ if and only if $\|M\|_\infty = 0$, or equivalently $M = 0$. Thus, the null space of the operator $\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega$ is the zero matrix. Hence, $\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega$ is a bijection, and thus invertible.

Next, we prove inequality (6.22). Let $(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)^{-1}M = M'$, or equivalently $M = (\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)M'$. Then, analogues to the derivation above, we have

$$\|M\|_\infty = \|(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)M'\|_\infty \geq (1 - 3\gamma)\|M'\|_\infty = (1 - 3\gamma)\|(\mathcal{I}_d - \mathcal{P}_{\bar{T}} \mathcal{P}_\Omega)^{-1}M\|_\infty,$$

which is inequality (6.22). \square

Appendix C: Bounding the diagonal deviation of the low-rank matrix estimator

We commented in the remark following Corollary 6.8 that the choice $c = 1$ in G_c is sufficient for proving a bound on $\|\tilde{\Sigma} - \Sigma\|_F^2$. On the other hand, exactly as commented in [48], and as

is apparent from Theorem 6.4, choosing $c > 1$ leads to a bound for $\mathcal{P}_{\tilde{T}^\perp} \tilde{\Theta}$, that is, the portion of $\tilde{\Theta}$ orthogonal to the tangent space \tilde{T} . As in [20], such a bound can be further exploited to control $\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)$, which in our case is the deviation of $\tilde{\Theta}$ from Θ^* on the diagonal. We first present a lemma toward the bound for $\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)$. The proof of the lemma is a straightforward modification of the proof of [20], Theorem 7; for completeness, we include it here. We employ the same notation as in Section 6.3, and we denote $E = \widehat{\Sigma} - \Sigma$ again.

Lemma C.1. *Let $r = \text{rank}(\tilde{\Theta})$. We have*

$$(1 - 3\gamma) \|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1 \leq \|\mathcal{P}_{\tilde{T}^\perp}(\tilde{\Theta} - \Theta^*)\|_* + 4r(\|E\|_2 + \mu). \tag{C.1}$$

Proof. Let $\tilde{\Delta}_\Theta = \tilde{\Theta} - \Theta^*$. The optimality of $\tilde{\Theta}$ for the convex program (3.13) implies that we can fix $\Psi \in \mu \partial \|\tilde{\Theta}\|_*$ such that equation (6.9) holds. Using $\nabla L(\tilde{\Theta}) = \tilde{\Theta}_o - \widehat{\Sigma}_o$, equation (6.9) is equivalent to

$$\tilde{\Delta}_\Theta = \mathcal{P}_\Omega \tilde{\Delta}_\Theta + E - \Psi. \tag{C.2}$$

Applying $\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}}$ on both sides of equation (C.2) gives

$$\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \tilde{\Delta}_\Theta = \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \mathcal{P}_\Omega \tilde{\Delta}_\Theta + \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} E - \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \Psi. \tag{C.3}$$

Then, using equation (C.3), we have

$$\begin{aligned} \mathcal{P}_\Omega \tilde{\Delta}_\Theta &= \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta + \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \tilde{\Delta}_\Theta \\ &= \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta + \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \mathcal{P}_\Omega \tilde{\Delta}_\Theta + \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} E - \mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \Psi. \end{aligned} \tag{C.4}$$

We apply $\|\cdot\|_1$ on both sides of equation (C.4). Note that, for any matrix $M \in \mathbb{R}^{d \times d}$, $\|\mathcal{P}_\Omega M\|_1 = \|\mathcal{P}_\Omega M\|_*$. In addition, inequality (6.21) implies that $\|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1 \leq 3\gamma \|\mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1$. Hence, we have

$$\begin{aligned} \|\mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1 &\leq \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta\|_1 + \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1 + \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} E\|_1 + \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \Psi\|_1 \\ &\leq \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta\|_* + 3\gamma \|\mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1 + \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} E\|_* + \|\mathcal{P}_\Omega \mathcal{P}_{\tilde{T}} \Psi\|_*. \end{aligned} \tag{C.5}$$

Note that, for any matrix $M \in \mathbb{R}^{d \times d}$, we have $\mathcal{P}_\Omega M = I_d \circ M$. By [19], Theorem 5.5.19, $\|I_d \circ M\|_* \leq \|M\|_*$. In addition, $\text{rank}(\mathcal{P}_{\tilde{T}} M) \leq 2r$, and so $\|\mathcal{P}_{\tilde{T}} M\|_* \leq 2r \|\mathcal{P}_{\tilde{T}} M\|_2 \leq 4r \|M\|_2$. Hence, from inequality (C.5), we further deduce

$$(1 - 3\gamma) \|\mathcal{P}_\Omega \tilde{\Delta}_\Theta\|_1 \leq \|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta\|_* + \|\mathcal{P}_{\tilde{T}} E\|_* + \|\mathcal{P}_{\tilde{T}} \Psi\|_* \leq \|\mathcal{P}_{\tilde{T}^\perp} \tilde{\Delta}_\Theta\|_* + 4r\|E\|_2 + 4r\|\Psi\|_2.$$

The corollary then follows by noting that $\|\Psi\|_2 \leq \mu$. □

We now state a concrete bound for $\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)$.

Theorem C.2. *Let μ and $\bar{\mu}$ be as in (3.15) and (3.16), respectively, and let*

$$\mu' = C \{ C_1 \max[\sqrt{\|T\|_2} f(n, d, \alpha), f^2(n, d, \alpha)] + C_2 f^2(n, d, \alpha) \}, \tag{C.6}$$

all with $0 < \alpha < 1/2$, $C_1 = \pi$, $C_2 = 3\pi^2/16 < 1.86$, and $C = 6$. We recall R as defined in (3.17). Then, with probability exceeding $1 - 2\alpha$, we have

$$\begin{aligned} & \|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1 \\ & \leq \min_{0 \leq r \leq R} \left\{ \frac{3}{2\mu'} \sum_{j: r < j \leq r^*} \lambda_j^2(\Theta^*) + \frac{3}{2} \sum_{j: r < j \leq r^*} \lambda_j(\Theta^*) + 19r\bar{\mu} \right\}. \end{aligned} \quad (\text{C.7})$$

Proof. We fix $c = 2$, and $\gamma' = 1/9$. Then inequality (6.33) holds with the substitution of γ_r by γ' . Let A be the event

$$A = \left\{ \left(\frac{1}{c} - \frac{\gamma'}{1 - 3\gamma'} \right)^{-1} \left(\frac{2\sqrt{\gamma'}}{1 - 3\gamma'} + 1 \right) \|E\|_2 \leq \mu' \leq \mu \leq \bar{\mu} \right\}. \quad (\text{C.8})$$

Hence, on the event A , both $\mu' \leq \mu \leq \bar{\mu}$, and inequality (6.34) with the substitution of γ_r by γ' , hold. Note that the multiplicative factor in front of $\|E\|_2$ on the right-hand side of (C.8) exactly equals $C = 6$ with our choices of c and γ' . Then, by Theorem 2.2 and our choices (C.6), (3.15) and (3.16) of μ' , μ and $\bar{\mu}$, we conclude that $\mathbb{P}(A) \geq 1 - \alpha - \alpha^2/4 > 1 - 2\alpha$, and for the rest of the proof we focus on the event A .

Note that Lemma C.1 provides a bound on $\|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1$ through the chosen $\tilde{\Theta}$ and the associated \tilde{T}^\perp . We fix an arbitrary $0 \leq r \leq R$, and choose $\tilde{\Theta} = \Theta_r^*$, which implies that $\gamma = \gamma_r$. Then

$$\begin{aligned} \mathcal{P}_{\tilde{T}^\perp}(\tilde{\Theta} - \Theta^*) &= \mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta} - \mathcal{P}_{T(\Theta_r^*)^\perp} \Theta^* \\ &= \mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta} - (\Theta^* - \Theta_r^*) \end{aligned}$$

and so

$$\|\mathcal{P}_{\tilde{T}^\perp}(\tilde{\Theta} - \Theta^*)\|_* \leq \|\mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta}\|_* + \sum_{j>r} \lambda_j(\Theta^*). \quad (\text{C.9})$$

Plugging inequality (C.9) into inequality (C.1) with the substitution of γ by γ_r yields

$$\begin{aligned} & \|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1 \\ & \leq \left(\frac{1}{1 - 3\gamma_r} \right) \left[\|\mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta}\|_* + \sum_{j>r} \lambda_j(\Theta^*) + 4r(\|E\|_2 + \mu) \right]. \end{aligned} \quad (\text{C.10})$$

As argued in the proof of Theorem 3.2, because inequalities (6.33) and (6.34) hold with the substitution of γ_r by γ' , we conclude that inequalities (6.33) and (6.34) hold in terms of γ_r . Hence, by Corollary 6.8, inequality (6.35) applies, and we have

$$\|\mathcal{P}_{T(\Theta_r^*)^\perp} \tilde{\Theta}\|_* \leq \frac{1}{\mu} \left[\sum_{j>r} \lambda_j^2(\Theta^*) + 8r\mu^2 \right]. \quad (\text{C.11})$$

Plugging inequality (C.11) into inequality (C.10), we have

$$\begin{aligned} \|\mathcal{P}_\Omega(\tilde{\Theta} - \Theta^*)\|_1 &\leq \left(\frac{1}{1-3\gamma_r}\right) \left\{ \frac{1}{\mu} \left[\sum_{j>r} \lambda_j^2(\Theta^*) + 8r\mu^2 \right] + \sum_{j>r} \lambda_j(\Theta^*) + 4r(\|E\|_2 + \mu) \right\} \\ &\leq \frac{3}{2} \left\{ \frac{1}{\mu} \sum_{j>r} \lambda_j^2(\Theta^*) + \sum_{j>r} \lambda_j(\Theta^*) + \frac{38}{3}r\mu \right\} \\ &\leq \frac{3}{2} \left\{ \frac{1}{\mu'} \sum_{j>r} \lambda_j^2(\Theta^*) + \sum_{j>r} \lambda_j(\Theta^*) + \frac{38}{3}r\bar{\mu} \right\}. \end{aligned} \quad (\text{C.12})$$

Here, the second inequality follows because $\gamma_r \leq 1/9$ and $\|E\|_2 \leq \mu/6$, and the last inequality follows because $\mu' \leq \mu \leq \bar{\mu}$. Then inequality (C.7) is obtained by minimizing inequality (C.12) over $0 \leq r \leq R$. \square

Acknowledgements

We are grateful to Han Liu for helpful discussions and for providing independent credit to our work. We thank the area Editor and the referees for their very constructive comments. This research is supported in part by NSF Grants DMS-10-07444 and DMS-13-10119.

References

- [1] Agarwal, A., Negahban, S. and Wainwright, M.J. (2012). Noisy matrix decomposition via convex relaxation: Optimal rates in high dimensions. *Ann. Statist.* **40** 1171–1197. [MR2985947](#)
- [2] Bickel, P.J. and Levina, E. (2008). Covariance regularization by thresholding. *Ann. Statist.* **36** 2577–2604. [MR2485008](#)
- [3] Bunea, F., She, Y. and Wegkamp, M.H. (2011). Optimal selection of reduced rank estimators of high-dimensional matrices. *Ann. Statist.* **39** 1282–1309. [MR2816355](#)
- [4] Bunea, F. and Xiao, L. (2015). On the sample covariance matrix estimator of reduced effective rank population matrices, with applications to fPCA. *Bernoulli* **21** 1200–1230. [MR3338661](#)
- [5] Cai, T.T., Zhang, C.-H. and Zhou, H.H. (2010). Optimal rates of convergence for covariance matrix estimation. *Ann. Statist.* **38** 2118–2144. [MR2676885](#)
- [6] Cai, T.T. and Zhou, H.H. (2012). Optimal rates of convergence for sparse covariance matrix estimation. *Ann. Statist.* **40** 2389–2420. [MR3097607](#)
- [7] Cambanis, S., Huang, S. and Simons, G. (1981). On the theory of elliptically contoured distributions. *J. Multivariate Anal.* **11** 368–385. [MR0629795](#)
- [8] Candès, E.J. and Recht, B. (2009). Exact matrix completion via convex optimization. *Found. Comput. Math.* **9** 717–772. [MR2565240](#)
- [9] Chandrasekaran, V., Parrilo, P.A. and Willsky, A.S. (2012). Latent variable graphical model selection via convex optimization. *Ann. Statist.* **40** 1935–1967. [MR3059067](#)
- [10] Chandrasekaran, V., Sanghavi, S., Parrilo, P.A. and Willsky, A.S. (2011). Rank-sparsity incoherence for matrix decomposition. *SIAM J. Optim.* **21** 572–596. [MR2817479](#)
- [11] Demarta, S. and McNeil, A.J. (2005). The t copula and related copulas. *Int. Stat. Rev.* **73** 111–129.

- [12] Eckart, C. and Young, G. (1936). The approximation of one matrix by another of lower rank. *Psychometrika* **1** 211–218.
- [13] Embrechts, P., Lindskog, F. and McNeil, A. (2003). Modelling dependence with copulas and applications to risk management. In *Handbook of Heavy Tailed Distributions in Finance* (S.T. Rachev, ed.) 329–384. Amsterdam: Elsevier.
- [14] Fang, H.-B., Fang, K.-T. and Kotz, S. (2002). The meta-elliptical distributions with given marginals. *J. Multivariate Anal.* **82** 1–16. [MR1918612](#)
- [15] Fazel, M. (2002). Matrix rank minimization with applications. Ph.D. thesis, Stanford Univ.
- [16] Friedman, J., Hastie, T. and Tibshirani, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9** 432–441.
- [17] Han, F. and Liu, H. (2015). Optimal rates of convergence for latent generalized correlation matrix estimation in transelliptical distribution. *Bernoulli*. To appear. Available at [arXiv:1305.6916](#).
- [18] Hoeffding, W. (1963). Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58** 13–30. [MR0144363](#)
- [19] Horn, R.A. and Johnson, C.R. (1991). *Topics in Matrix Analysis*. Cambridge: Cambridge Univ. Press. [MR1091716](#)
- [20] Hsu, D., Kakade, S.M. and Zhang, T. (2011). Robust matrix decomposition with sparse corruptions. *IEEE Trans. Inform. Theory* **57** 7221–7234. [MR2883652](#)
- [21] Hult, H. and Lindskog, F. (2002). Multivariate extremes, aggregation and dependence in elliptical distributions. *Adv. in Appl. Probab.* **34** 587–608. [MR1929599](#)
- [22] Kendall, M.G. and Gibbons, J.D. (1990). *Rank Correlation Methods*, 5th ed. London: Edward Arnold. [MR1079065](#)
- [23] Klüppelberg, C. and Kuhn, G. (2009). Copula structure analysis. *J. R. Stat. Soc. Ser. B. Stat. Methodol.* **71** 737–753. [MR2749917](#)
- [24] Klüppelberg, C., Kuhn, G. and Peng, L. (2008). Semi-parametric models for the multivariate tail dependence function—The asymptotically dependent case. *Scand. J. Stat.* **35** 701–718. [MR2468871](#)
- [25] Koltchinskii, V., Lounici, K. and Tsybakov, A.B. (2011). Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Statist.* **39** 2302–2329. [MR2906869](#)
- [26] Kruskal, W.H. (1958). Ordinal measures of association. *J. Amer. Statist. Assoc.* **53** 814–861. [MR0100941](#)
- [27] Lindskog, F., McNeil, A. and Schmock, U. (2003). Kendall’s tau for elliptical distributions. In *Credit Risk: Measurement, Evaluation and Management, Contributions to Economics* (G. Bol, G. Nakhaeizadeh, S.T. Rachev, T. Ridder and K.-H. Vollmer, eds.) 149–156. Heidelberg: Physica-Verlag.
- [28] Liu, H., Han, F., Yuan, M., Lafferty, J. and Wasserman, L. (2012). High-dimensional semiparametric Gaussian copula graphical models. *Ann. Statist.* **40** 2293–2326. [MR3059084](#)
- [29] Liu, H., Han, F. and Zhang, C.-H. (2012). Transelliptical graphical models. In *Adv. Neural Inf. Process. Syst.* (P. Bartlett, F.C.N. Pereira, C.J.C. Burges, L. Bottou and K.Q. Weinberger, eds.) **25** 809–817. Neural Information Processing Systems Foundation.
- [30] Liu, H., Lafferty, J. and Wasserman, L. (2009). The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *J. Mach. Learn. Res.* **10** 2295–2328. [MR2563983](#)
- [31] Lounici, K. (2014). High-dimensional covariance matrix estimation with missing observations. *Bernoulli* **20** 1029–1058. [MR3217437](#)
- [32] Luo, X. (2013). Recovering model structures from large low rank and sparse covariance matrix estimation. Preprint. Available at [arXiv:1111.1133](#).
- [33] Mitra, R. and Zhang, C.-H. (2014). Multivariate analysis of nonparametric estimates of large correlation matrices. Preprint. Available at [arXiv:1403.6195](#).

- [34] Negahban, S. and Wainwright, M.J. (2011). Estimation of (near) low-rank matrices with noise and high-dimensional scaling. *Ann. Statist.* **39** 1069–1097. [MR2816348](#)
- [35] Petz, D. (1994). A survey of certain trace inequalities. In *Functional Analysis and Operator Theory (Warsaw, 1992)*. *Banach Center Publ.* **30** 287–298. Warsaw: Polish Acad. Sci. [MR1285615](#)
- [36] Qi, H. and Sun, D. (2006). A quadratically convergent Newton method for computing the nearest correlation matrix. *SIAM J. Matrix Anal. Appl.* **28** 360–385. [MR2255334](#)
- [37] Rohde, A. and Tsybakov, A.B. (2011). Estimation of high-dimensional low-rank matrices. *Ann. Statist.* **39** 887–930. [MR2816342](#)
- [38] Saunderson, J., Chandrasekaran, V., Parrilo, P.A. and Willsky, A.S. (2012). Diagonal and low-rank matrix decompositions, correlation matrices, and ellipsoid fitting. *SIAM J. Matrix Anal. Appl.* **33** 1395–1416. [MR3028972](#)
- [39] Schmidt, E. (1907). Zur Theorie der linearen und nichtlinearen Integralgleichungen. *Math. Ann.* **63** 433–476. [MR1511415](#)
- [40] Sklar, A. (1996). Random variables, distribution functions, and copulas—A personal look backward and forward. In *Distributions with Fixed Marginals and Related Topics (Seattle, WA, 1993)* (L. Rüschendorf, B. Schweizer and M. D. Taylor, eds.). *Institute of Mathematical Statistics Lecture Notes—Monograph Series* **28** 1–14. Hayward, CA: IMS. [MR1485519](#)
- [41] Tropp, J.A. (2012). User-friendly tail bounds for sums of random matrices. *Found. Comput. Math.* **12** 389–434. [MR2946459](#)
- [42] Tropp, J.A. (2014). An introduction to matrix concentration inequalities. Technical report, California Institute of Technology.
- [43] Vershynin, R. (2012). Introduction to the nonasymptotic analysis of random matrices. In *Compressed Sensing* (Y. Eldar and G. Kutyniok, eds.) *Compressed Sensing, Theory and Application* 210–268. Cambridge: Cambridge Univ. Press. [MR2963170](#)
- [44] Watson, G.A. (1992). Characterization of the subdifferential of some matrix norms. *Linear Algebra Appl.* **170** 33–45. [MR1160950](#)
- [45] Xue, L., Ma, S. and Zou, H. (2012). Positive-definite ℓ_1 -penalized estimation of large covariance matrices. *J. Amer. Statist. Assoc.* **107** 1480–1491. [MR3036409](#)
- [46] Xue, L. and Zou, H. (2012). Regularized rank-based estimation of high-dimensional nonparanormal graphical models. *Ann. Statist.* **40** 2541–2571. [MR3097612](#)
- [47] Yuan, M. (2012). Comment: “Minimax estimation of large covariance matrices under ℓ_1 -norm” [[MR3027084](#)]. *Statist. Sinica* **22** 1373–1375. [MR3027090](#)
- [48] Zhang, C.-H. and Zhang, T. (2012). A general framework of dual certificate analysis for structured sparse recovery problems. Technical report, Rutgers Univ.

Received January 2014 and revised September 2014