

On estimation of the diagonal elements of a sparse precision matrix

Samuel Balmand

Université Paris-Est, IGN, MATIS, ENSG, F-77455, Marne-la-Vallée, France
e-mail: samuel.balmand@ensg.eu

and

Arnak S. Dalalyan

ENSAE ParisTech & CREST, 3 av. P. Larousse, 92245 Malakoff, France
e-mail: arnak.dalalyan@ensae.fr

Abstract: In this paper, we present several estimators of the diagonal elements of the inverse of the covariance matrix, called precision matrix, of a sample of independent and identically distributed random vectors. The main focus is on the case of high dimensional vectors having a sparse precision matrix. It is now well understood that when the underlying distribution is Gaussian, the columns of the precision matrix can be estimated independently from one another by solving linear regression problems under sparsity constraints. This approach leads to a computationally efficient strategy for estimating the precision matrix that starts by estimating the regression vectors, then estimates the diagonal entries of the precision matrix and, in a final step, combines these estimators for getting estimators of the off-diagonal entries. While the step of estimating the regression vector has been intensively studied over the past decade, the problem of deriving statistically accurate estimators of the diagonal entries has received much less attention. The goal of the present paper is to fill this gap by presenting four estimators—that seem the most natural ones—of the diagonal entries of the precision matrix and then performing a comprehensive empirical evaluation of these estimators. The estimators under consideration are the residual variance, the relaxed maximum likelihood, the symmetry-enforced maximum likelihood and the penalized maximum likelihood. We show, both theoretically and empirically, that when the aforementioned regression vectors are estimated without error, the symmetry-enforced maximum likelihood estimator has the smallest estimation error. However, in a more realistic setting when the regression vector is estimated by a sparsity-favoring computationally efficient method, the qualities of the estimators become relatively comparable with a slight advantage for the residual variance estimator.

MSC 2010 subject classifications: Primary 62H12.

Keywords and phrases: Precision matrix, sparse recovery, penalized likelihood.

Received June 2015.

1. Introduction

We consider the problem of precision matrix estimation that has been extensively studied in recent years partly because of its tight relation with the graphical models. More precisely, assuming that we observe p features on n individuals, an interesting object to display is the graph of associations between the features, especially when the number of features is large. The associations may be of different type: linear correlations, partial correlations, measures of independence and so on. A measure of association between the features, which is particularly relevant for Gaussian (Lauritzen, 1996) and, more generally, non parnormal distributions (Liu et al., 2009; Lafferty et al., 2012) is the partial correlation. This leads to a Gaussian graphical model in which two nodes are connected by an edge if the partial correlation between the features corresponding to these two nodes is nonzero, which is equivalent to the nonzeroness of the corresponding entry of the precision matrix (Lauritzen, 1996, Proposition 5.2). The graph constructed in such a way relies on the population precision matrix, which is not available in practice. Therefore, an important statistical problem is to infer this graph from n iid observations of the p -dimensional feature-vector. In view of the aforementioned connection with the precision matrix, the estimated graph may be deduced from the estimated precision matrix by comparing its entries with a suitably chosen threshold.

Another important problem for which the precision matrix estimation is relevant¹ is the linear (Fisher, 1936) or quadratic discriminant analysis (Anderson, 2003). Indeed, the decision boundary in the binary or multi-class classification problem—under the assumption that the conditional distributions of the features given the class are Gaussian—is defined in terms of the precision matrix. In order to infer this decision boundary from data, it is therefore relevant to start with estimating the precision matrix. The simplest way of estimating the latter is by inverting the sample covariance matrix or, if the inverse does not exist, by computing the pseudo-inverse of the sample covariance matrix. However, when the dimension p is such that the number of unknown parameters $p(p+1)$ is comparable to or larger than the sample-size n , the (pseudo-)inversion of the sample covariance matrix leads to very poor results. To circumvent this shortcoming, additional assumptions on the precision matrix should be imposed which should preferably be realistic, interpretable and lead to statistically and computationally efficient estimation procedures. The sparsity of the precision matrix offers a convenient setting in which these criteria are met.

To present in a more concrete fashion the content of the present work, let \mathbf{X} be a $n \times p$ random matrix representing the values of p variables observed on n individuals. Assume that the rows of the matrix \mathbf{X} are independent and Gaussian with mean $\boldsymbol{\mu}^*$ and covariance matrix $\boldsymbol{\Sigma}^*$. The inverse of $\boldsymbol{\Sigma}^*$, called the precision matrix and denoted by $\boldsymbol{\Omega}^* = (\omega_{ij}^*)$, is an object of central interest

¹In the case of linear discriminant analysis for binary classification, a simpler approach consisting in replacing the sparsity of the precision matrix by the sparsity of the product of the latter with the difference of the class means has been proposed and studied by Cai and Liu (2011).

since—as mentioned earlier—it encodes the conditional dependencies between pairs of variables given the values of all the other variables. Based on the precision matrix, the graph \mathcal{G}^* of relationships between the p variables is constructed as follows: each node of the graph represents a variable and two nodes i and j are connected by an edge if and only if $\omega_{ij}^* \neq 0$. Estimating this graph from a sample of size n represented by the rows of \mathbf{X} is a challenging statistical problem that has attracted a lot of attention in the past decade. In a frequently encountered situation of the dimension p comparable to or even larger than n , a commonly used assumption is the sparsity of the graph \mathcal{G}^* . Namely, it is assumed that the maximal degree of the nodes of \mathcal{G}^* is much smaller than p (see, for instance, Meinshausen and Bühlmann (2006); Yuan and Lin (2007) for early references).

Most approaches of estimating sparse precision matrices that gained popularity in recent years rely on weighted ℓ_1 -penalization of the off-diagonal elements of the precision matrix; recent contributions on the statistical aspects of this approach can be found in Yuan (2010); Cai et al. (2011); Sun and Zhang (2013); Cai et al. (2016) and the references therein. The rationale behind this approach is that the weighted ℓ_1 -penalty can be viewed as a convexified version of the ℓ_0 -penalty, the latter being understood as the number of nonzero elements. The convexity of the penalty in conjunction with the convexity of the data fidelity term leads to estimators that can be efficiently computed by convex programming (Friedman et al., 2008; Banerjee et al., 2008).

To further improve the computational complexity, it is possible to split the problem of estimating p^2 entries of the precision matrix into p independent problems of estimating the p -dimensional columns of it (Meinshausen and Bühlmann, 2006). To this end, the matrix $\mathbf{\Omega}^*$ is written as $\mathbf{B}^* \mathbf{D}^*$, where \mathbf{D}^* is a diagonal matrix while \mathbf{B}^* is a $p \times p$ matrix with all diagonal entries equal to one. Each column of the matrix \mathbf{B}^* can be estimated by regressing one column of the data matrix \mathbf{X} on all the remaining columns. In the context of high dimensionality and sparse precision matrix, this can be performed by sparsity favoring methods (Bühlmann and van de Geer, 2011) such as the Lasso (Tibshirani, 1996), the Dantzig selector (Candes and Tao, 2007), the square-root Lasso (Belloni et al., 2011), etc. A crucial observation at this stage is that the sparsity patterns, *i.e.*, the locations of nonzero entries, of the matrices \mathbf{B}^* and $\mathbf{\Omega}^*$ coincide. In particular, the degree of the j -th node in the graph \mathcal{G}^* is equal to the number of nonzero entries of the j -th column of \mathbf{B}^* , for every $j = 1, \dots, p$.

Once the columns of \mathbf{B}^* successfully estimated, one needs to estimate the diagonal matrix \mathbf{D}^* , the diagonal entries of which coincide with those of the precision matrix $\mathbf{\Omega}^*$. This step is necessary for recovering the precision matrix (both diagonal and nondiagonal entries) but it is also important for constructing the graph² of conditional dependencies. Of course, the latter can be estimated by thresholding the entries of the estimator of \mathbf{B}^* without resorting to an estimator of \mathbf{D}^* , but the choice of the threshold is in this case a difficult issue deprived of clear statistical interpretation. In contrast with this, if along with an estimator of \mathbf{B}^* , an estimator of \mathbf{D}^* is available, then one may straightforwardly estimate

²We put an emphasize on this last point since we did not find it in the literature.

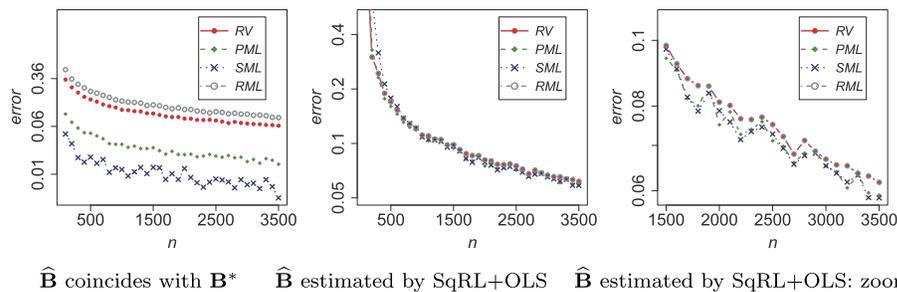


FIG 1. The average ℓ_2 -error (computed from 50 independent trials) of the four estimators considered in this work as a function of the sample size. The plots concern Model 2 described in Section 4.1 and dimension $p = 60$. One can observe, in particular, that when \mathbf{B}^* is estimated without error (left panel), the estimators SML and PML improve on the residual variance and relaxed maximum likelihood estimators.

the partial correlations and threshold them to infer the graph of conditional dependencies. In this case, the threshold has a more clear statistical meaning since the partial correlations are in absolute value bounded by one.

It follows from the above discussion that the problem of estimating the matrix \mathbf{D}^* built from the diagonal entries of the precision matrix is an important ingredient of the estimation of the precision matrix and the graph of conditional dependencies between the features. The purpose of the present work is to propose several natural estimators of \mathbf{D}^* and to study their statistical properties, essentially from an empirical point of view. Combining standard arguments, we present four estimators, termed residual variance (RV), relaxed maximum likelihood (RML), symmetry-enforced maximum likelihood (SML) and penalized maximum likelihood (PML). The first one, residual variance, is the most commonly used estimator when the matrix \mathbf{B}^* is estimated column-wise by a sparse linear regression approach briefly mentioned in the foregoing discussion. The other three methods considered in this paper are based on the principle of likelihood maximization under various approaches for handling the prior information. In order to give the reader a foretaste of the content of next sections, we present in Figure 1 the accuracy of the four methods of estimating the diagonal elements of the precision matrix on a synthetic data-set. More details are given in Section 4.1.

2. Notation and preliminaries on precision matrix estimation

This section introduces notation used throughout the paper and presents some preliminary material on sparse precision matrix estimation.

2.1. Notation

For an unknown parameter θ we note θ^* its true value. As usual, $\mathcal{N}_p(\boldsymbol{\mu}^*, \boldsymbol{\Sigma}^*)$ is the Gaussian distribution in \mathbb{R}^p with mean $\boldsymbol{\mu}^*$ and covariance matrix $\boldsymbol{\Sigma}^*$.

The expectation of a random vector \mathbf{X} is denoted by $\mathbf{E}(\mathbf{X})$ and its covariance matrix by $\mathbf{V}(\mathbf{X})$. We denote by $\mathbf{1}_n$ the vector from \mathbb{R}^n with all the entries equal to 1 and by \mathbf{I}_n the $n \times n$ identity matrix. We write $\mathbb{1}$ for the indicator function, which is equal to 1 if the considered condition is satisfied and 0 otherwise. The smallest integer greater than or equal to $x \in \mathbb{R}$ is denoted by $\lceil x \rceil$. In what follows, $[p] := \{1, \dots, p\}$ is the set of positive integers from 1 to p . For $i \in [p]$, the complement of the singleton $\{i\}$ in $[p]$ is denoted by i^c . For a vector $\mathbf{v} \in \mathbb{R}^p$, \mathbf{D}_v stands for the $p \times p$ diagonal matrix satisfying $(\mathbf{D}_v)_j = v_j$ for every $j \in [p]$. The matrix build keeping only the diagonal of a square matrix \mathbf{M} is denoted by $\text{diag}(\mathbf{M})$.

The transpose of the matrix \mathbf{M} is denoted by \mathbf{M}^\top . If this matrix is square, we note $|\mathbf{M}|$ its determinant. For a $n \times p$ matrix \mathbf{M} , the vector of the elements of the k th row (resp. the j th column) whose indices are given by the subset J of $[p]$ (resp. K of $[n]$) is denoted by $\mathbf{M}_{k,J}$ (resp. $\mathbf{M}_{K,j}$). In particular, the vector made of all the elements of the j th column of the matrix \mathbf{M} at the exception of the element of the k th row is given by $\mathbf{M}_{k^c,j}$. Moreover, the whole k th row (resp. j th column) of \mathbf{M} is denoted by $\mathbf{M}_{k,\bullet}$ (resp. $\mathbf{M}_{\bullet,j}$). We use the following notation for the (pseudo-)norms of matrices: if $q_1, q_2 > 0$, then

$$\|\mathbf{M}\|_{q_1, q_2} = \left\{ \sum_{i=1}^n \|\mathbf{M}_{i,\bullet}\|_{q_1}^{q_2} \right\}^{1/q_2}.$$

With this notation, $\|\mathbf{M}\|_{2,2}$ and $\|\mathbf{M}\|_{1,1}$ are the Frobenius and the element-wise ℓ_1 -norm of \mathbf{M} , respectively. The sample covariance matrix of the data points $\{\mathbf{X}_{k,\bullet}\}_{k \in [n]}$ is defined by

$$\mathbf{S}_n = \frac{1}{n} \sum_{k=1}^n (\mathbf{X}_{k,\bullet}^\top - \hat{\boldsymbol{\mu}})(\mathbf{X}_{k,\bullet}^\top - \hat{\boldsymbol{\mu}})^\top = \frac{1}{n} (\mathbf{X} - \mathbf{1}_n \hat{\boldsymbol{\mu}}^\top)^\top (\mathbf{X} - \mathbf{1}_n \hat{\boldsymbol{\mu}}^\top),$$

where $\hat{\boldsymbol{\mu}}$ is either the sample mean $\frac{1}{n}(\mathbf{1}_n^\top \mathbf{X})^\top$ (when the mean $\boldsymbol{\mu}^*$ is unknown) or the theoretical mean $\boldsymbol{\mu}^*$ (when it is considered as known).

2.2. Preliminaries

Throughout the paper we will present estimators of the diagonal elements of the precision matrix in the case of a general multidimensional Gaussian distribution, but in all theoretical developments we will assume that the marginals of \mathbf{X} are standard Gaussian distributions, *i.e.*, $\boldsymbol{\mu}^* = 0$ and $\boldsymbol{\Sigma}_{jj}^* = 1$ for every $j \in [p]$. This assumption is reasonable, since we are concerned with the problems in which the sample size is large enough to consistently estimate the individual means and the individual variances of the variables. So, one can always center the variables by the sample mean and divide by the sample standard deviation to get close to the assumption³ that random variables $\mathbf{X}_{1,j}, \dots, \mathbf{X}_{n,j}$ are i.i.d. $\mathcal{N}(0, 1)$ for every j .

³Unless expressly stated otherwise, in the whole article, $1 \leq i, j \leq p$ and $1 \leq k \leq n$.

Let us recall that the precision matrix is closely related to the problem of regression of one feature on all the others. Indeed, there exists a $p \times p$ matrix \mathbf{B}^* and two vectors $\mathbf{c}^*, \boldsymbol{\phi}^* \in \mathbb{R}^p$ such that

$$\mathbf{X}_{\bullet,j} = c_j^* \mathbf{1}_n - \mathbf{X}_{\bullet,j^c} \mathbf{B}_{j^c,j}^* + \phi_j^{*1/2} \boldsymbol{\xi}_j, \quad (1)$$

where $\boldsymbol{\xi}_j$ is drawn from $\mathcal{N}_n(0, \mathbf{I}_n)$ and is independent of \mathbf{X}_{\bullet,j^c} . According to the theorem on normal correlations (Marsaglia, 1964), the regression coefficients $\mathbf{B}_{j^c,j}^* \in \mathbb{R}^{p-1}$ and the variance $\phi_j^* \in \mathbb{R}$ of residuals can be expressed in terms of the elements of the precision matrix $\boldsymbol{\Omega}^*$ as follows:

$$\mathbf{B}_{ij}^* = \omega_{ij}^* / \omega_{jj}^*, \quad \phi_j^* = 1 / \omega_{jj}^*, \quad (\text{NC})$$

whereas $c_j^* = \mu_j^* + (\boldsymbol{\mu}_{j^c}^*)^\top \mathbf{B}_{j^c,j}^* = (\boldsymbol{\mu}^*)^\top \mathbf{B}_{\bullet,j}^*$. If we assume that $\boldsymbol{\mu}^* = 0$ then $c_j^* = 0$ for any j . With these notation, the precision matrix can be written as $\boldsymbol{\Omega}^* = \mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}^*}^{-1}$.

Several state-of-the-art methods for estimating sparse precision matrices proceed in two steps (Meinshausen and Bühlmann, 2006; Cai et al., 2011; Liu and Wang, 2012; Sun and Zhang, 2013). The first step consists in estimating the matrix \mathbf{B}^* and the vector $\boldsymbol{\phi}^*$ by solving the sparse linear regression problems (1) for each j , while in the second step an estimator of the matrix $\boldsymbol{\Omega}^*$ is inferred from the estimators of \mathbf{B}^* and $\boldsymbol{\phi}^*$ using relations (NC). The goal of the present work is to explore both theoretically and empirically different possible strategies for this second step.

The square-root Lasso is perhaps the method of estimating the matrix \mathbf{B}^* that offers the best trade-off between the computational and the statistical complexities. It can be redefined as follows: scaled Lasso estimates the matrix \mathbf{B}^* by solving the convex optimization problem

$$\widehat{\mathbf{B}} = \arg \min_{\substack{\mathbf{B} \in \mathbb{R}^{p \times p} \\ \mathbf{B}_{jj} = 1}} \min_{\mathbf{c} \in \mathbb{R}^p} \left\{ \|\mathbf{X}\mathbf{B} - \mathbf{1}_n \mathbf{c}^\top\|_{2,1} + \lambda \|\mathbf{B}\|_{1,1} \right\}, \quad (2)$$

where the first min is over all matrices \mathbf{B} having all their diagonal entries equal to 1. The tuning parameter $\lambda > 0$ corresponds to the penalty level. The purpose of the penalization is indeed to get a precision matrix estimate which fits the sparsity assumption. As the penalty of a matrix \mathbf{B} is its $\|\cdot\|_{1,1}$ norm, the resulting precision matrix estimator is expected to be sparse in the sense that its overall number of non-zero elements should be small. In addition, one can check that computing a solution to problem (2) is equivalent to computing each column of $\widehat{\mathbf{B}}$ separately (and independently) by solving the optimization problem

$$\widehat{\mathbf{B}}_{\bullet,j} = \arg \min_{\substack{\boldsymbol{\beta} \in \mathbb{R}^p \\ \beta_j = 1}} \min_{c_j \in \mathbb{R}} \left\{ \|\mathbf{X}\boldsymbol{\beta}_j - c_j \mathbf{1}_n\|_2 + \lambda \|\boldsymbol{\beta}_j\|_1 \right\}, \quad j \in [p]. \quad (3)$$

In addition to being efficiently computable even for large p , this estimator has the following appealing property that makes it preferable, for instance, to the

column-wise Lasso (Meinshausen and Bühlmann, 2006) and the CLIME (Cai et al., 2011). The choice of the parameter λ in (2-3) is scale free: it can be chosen independently of the noise variance in linear regression (1). This fact has been first established by Belloni et al. (2011) and then further investigated in (Sun and Zhang, 2012; Belloni et al., 2014a). In the context of precision matrix estimation, this method has been explored⁴ by Sun and Zhang (2013).

3. Four estimators of ϕ^*

As mentioned earlier, the aim of this work is to compare different estimators of the vector ϕ^* based on an initial estimator of the matrix \mathbf{B}^* . Clearly, the error of the estimation of \mathbf{B}^* impacts the error of the estimation of ϕ^* and, therefore, the latter is not easy to assess in full generality. In order to gain some insight on the behavior of various natural estimators, in theoretical results we will consider the ideal situation where the matrix \mathbf{B}^* is estimated without error.

3.1. Residual variance estimator

In view of the regression equation (1), a standard and natural method—used, in particular, by the square-root Lasso of Sun and Zhang (2013)—to deduce estimators $\hat{\phi}$ and $\hat{\Omega}$ from an estimator $\hat{\mathbf{B}}$ is to set

$$\hat{\phi}_j = \frac{1}{n} \|(\mathbf{I}_n - n^{-1} \mathbf{1}_n \mathbf{1}_n^\top) \mathbf{X} \hat{\mathbf{B}}_{\bullet,j}\|_2; \quad \hat{\Omega} = \hat{\mathbf{B}} \cdot \mathbf{D}_{\hat{\phi}}^{-1}. \quad (4)$$

Note that the matrix $(\mathbf{I}_n - n^{-1} \mathbf{1}_n \mathbf{1}_n^\top)$ present in this expression is the orthogonal projector in \mathbb{R}^n onto the orthogonal complement of the linear subspace $\text{Span}(\mathbf{1}_n)$ of all constant vectors. The multiplication by this matrix annihilates the intercept c_j^* in (1) and is a standard way of reducing the affine regression to the linear regression. In what follows, we refer to $\hat{\phi}$ defined by (4) as the residual variance estimator and denote it by $\hat{\phi}^{\text{RV}}$. Using the sample covariance matrix \mathbf{S}_n , the residual variance estimator of ϕ^* can be written as

$$\hat{\phi}_j^{\text{RV}} = \hat{\mathbf{B}}_{\bullet,j}^\top \mathbf{S}_n \hat{\mathbf{B}}_{\bullet,j}.$$

Note also that if we consider the linear regression model (1) conditionally to \mathbf{X}_{\bullet,j^c} , then the residual variance estimator of ϕ_j^* coincides with the maximum likelihood estimator.

Proposition 1. *If $\hat{\mathbf{B}}_{\bullet,j}$ estimates $\mathbf{B}_{\bullet,j}^*$ without error, then the residual variance estimator of ϕ_j^* has a quadratic risk equal to $\frac{2}{n} \phi_j^{*2}$, that is*

$$\mathbf{E}[(\hat{\phi}_j^{\text{RV}} - \phi_j^*)^2] = \frac{2\phi_j^{*2}}{n}.$$

⁴Although Sun and Zhang (2012, 2013) refer to this method as the scaled Lasso, we prefer to use the original term square-root Lasso coined by Belloni et al. (2011) in order to avoid any possible confusion with the earlier method of Städler et al. (2010a,b), for which the term “scaled Lasso” has been already employed.

Furthermore, for every $t > 0$, the following bound on the tails of the maximal error holds true:

$$\mathbf{P} \left(\max_{j \in [p]} \frac{|\widehat{\phi}_j^{\text{RV}} - \phi_j^*|}{\phi_j^*} > 2 \left(\frac{t + \log p}{n} \right)^{1/2} + \frac{2(t + \log p)}{n} \right) \leq 2e^{-t}.$$

Proof. Using equation (1) and the assumption $\widehat{\mathbf{B}}_{\bullet,j} = \mathbf{B}_{\bullet,j}^*$, we get

$$\widehat{\phi}_j^{\text{RV}} = \frac{1}{n} \|\mathbf{X} \mathbf{B}_{\bullet,j}^*\|_2^2 = \frac{\phi_j^*}{n} \|\boldsymbol{\xi}_j\|_2^2. \tag{5}$$

Since $\boldsymbol{\xi}_j$ is a standard Gaussian vector, the random variable $\zeta = \|\boldsymbol{\xi}_j\|_2^2$ is drawn from a χ_n^2 distribution. This implies that $\mathbf{E}(\zeta) = n$ and $\mathbf{V}(\zeta) = 2n$. Therefore,

$$\mathbf{E}[(\widehat{\phi}_j^{\text{RV}} - \phi_j^*)^2] = \mathbf{E} \left[\left(\frac{\phi_j^* \zeta}{n} - \phi_j^* \right)^2 \right] = \frac{\phi_j^{*2}}{n^2} (\mathbf{V}(\zeta) + (\mathbf{E}(\zeta) - n)^2) = \frac{2\phi_j^{*2}}{n}.$$

This completes the proof of the first claim. To prove the second claim, we set $z = t + \log p$ and use the union bound to get

$$\begin{aligned} \mathbf{P} \left(\max_{j \in [p]} \frac{|\widehat{\phi}_j^{\text{RV}} - \phi_j^*|}{\phi_j^*} > 2\sqrt{\frac{z}{n}} + \frac{2z}{n} \right) &\leq p \max_{j \in [p]} \mathbf{P} \left(\frac{|\widehat{\phi}_j^{\text{RV}} - \phi_j^*|}{\phi_j^*} > 2\sqrt{\frac{z}{n}} + \frac{2z}{n} \right) \\ &= p \mathbf{P}(|\zeta - n| > 2\sqrt{zn} + 2z). \end{aligned}$$

The second claim follows from the tail bound of the χ^2 distribution established, for instance, in (Laurent and Massart, 2000, Lemma 1). \square

Note that in this result, the case of known means μ_j^* is considered. The case of unknown μ_j can be handled similarly, the estimation bias is then ϕ_j^*/n and the resulting mean squared error is $(2n - 1)\phi_j^{*2}/n^2$. One may observe that, as expected, the rate of convergence of the quadratic risk is the usual parametric rate $1/n$ and that the asymptotic variance is $2\phi_j^{*2}$.

3.2. Relaxed maximum likelihood estimator

One could expect that the global maximum likelihood estimator of $\boldsymbol{\phi}^*$ would be better than the maximum of the conditional likelihood, since it is well known that under proper regularity conditions, the quadratic risk of the maximum likelihood estimator is the smallest, at least asymptotically. Since the vectors $\mathbf{X}_{k,\bullet} \sim \mathcal{N}_p(\boldsymbol{\mu}^*, \boldsymbol{\Omega}^{*-1})$ are independent, the log-likelihood is given by (up to irrelevant additive terms independent of the unknown parameters $\boldsymbol{\mu}^*$ and $\boldsymbol{\Omega}^*$)

$$\mathcal{L}(\mathbf{X}|\boldsymbol{\mu}, \boldsymbol{\Omega}) = \frac{n}{2} \log |\boldsymbol{\Omega}| - \frac{1}{2} \sum_{k=1}^n (\mathbf{X}_{k,\bullet} - \boldsymbol{\mu}^\top) \boldsymbol{\Omega} (\mathbf{X}_{k,\bullet} - \boldsymbol{\mu}^\top)^\top. \tag{6}$$

Maximizing the log-likelihood with respect to $\boldsymbol{\mu} \in \mathbb{R}^p$ leads to

$$\max_{\boldsymbol{\mu} \in \mathbb{R}^p} \mathcal{L}(\mathbf{X}|\boldsymbol{\mu}, \boldsymbol{\Omega}) = \frac{n}{2} \left(\log |\boldsymbol{\Omega}| - \text{trace}[\mathbf{S}_n \boldsymbol{\Omega}] \right). \tag{7}$$

Recall now that in view of (NC), we have $\boldsymbol{\Omega}^* = \mathbf{B}^* \mathbf{D}_{\phi^*}^{-1}$. Therefore, the profiled log-likelihood (w.r.t. $\boldsymbol{\mu}$) of \mathbf{X} given the parameters \mathbf{B} and $\boldsymbol{\phi}$ is

$$\max_{\boldsymbol{\mu} \in \mathbb{R}^p} \mathcal{L}(\mathbf{X} | \boldsymbol{\mu}, \mathbf{B}, \boldsymbol{\phi}) = \frac{n}{2} \left(\log |\mathbf{B}| - \sum_{j=1}^p \{ \log(\phi_j) + (\mathbf{S}_n \mathbf{B})_{jj} \phi_j^{-1} \} \right). \quad (8)$$

For a given \mathbf{B} , this profiled log-likelihood is a decomposable function of $\boldsymbol{\phi}$ and, therefore, can be easily maximized with respect to $\boldsymbol{\phi}$. This leads to

$$\arg \max_{\boldsymbol{\phi} \in \mathbb{R}_+^p} \max_{\boldsymbol{\mu} \in \mathbb{R}^p} \mathcal{L}(\mathbf{X} | \boldsymbol{\mu}, \mathbf{B}, \boldsymbol{\phi}) = ((\mathbf{S}_n \mathbf{B})_{jj} \vee 0)_{j \in [p]}. \quad (9)$$

Thus, when an estimator $\widehat{\mathbf{B}}$ of \mathbf{B}^* is available, one possible approach for estimating $\boldsymbol{\phi}^*$ is to set

$$\widehat{\phi}_j^{\text{RML}} = (\mathbf{S}_n \widehat{\mathbf{B}})_{jj} \vee 0, \quad j \in [p]. \quad (10)$$

We call this estimator relaxed maximum likelihood (RML) estimator. It will be clear a little bit later why it is called relaxed. The analysis of the risk of the RML estimator is more involved than that of the RV estimator considered in the previous section. This is due to the truncation at the level 0. For this reason, the next result does not provide the precise value of the risk, but just an inequality which is sufficient for our purposes.

Proposition 2. *If $\widehat{\mathbf{B}}$ estimates $\mathbf{B}_{\bullet,j}^*$ without error, then the risk of the RML estimator of ϕ_j^* satisfies $\mathbf{E}[(\widehat{\phi}_j^{\text{RML}} - \phi_j^*)^2] \geq \frac{1}{n}(\phi_j^{*2} + \phi_j^* \boldsymbol{\Sigma}_{jj}^* - O(n^{-1/2}))$.*

Before providing the proof of this result, let us present a brief discussion. Note that in view of (1), $\boldsymbol{\Sigma}_{jj}^*$ is always not smaller than ϕ_j^* . Furthermore, $\boldsymbol{\Sigma}_{jj}^* > \phi_j^*$ if $\mathbf{B}_{j^c,j}^*$ has at least one nonzero entry. Therefore, the last proposition, combined with Proposition 1, establishes that the residual variance estimator has an asymptotic variance which is smaller (and, in many cases, strictly smaller) than the asymptotic variance of the maximum likelihood estimator. At a first sight, this is very surprising and seems to be in contradiction with the well established theory (Le Cam and Yang, 2000; Ibragimov and Has'minskiĭ, 1981) of asymptotic efficiency of the maximum likelihood estimator for regular models. Our explanation of this inefficiency of $\widehat{\phi}_j^{\text{RML}}$ is that it is not really the maximum likelihood estimator. It maximizes the likelihood, certainly, but not over the correct set of parameters. Indeed, when we defined the RML we neglected an important property of the vector $\boldsymbol{\phi}^*$: the fact that $\mathbf{B}^* \mathbf{D}_{\phi^*}^{-1} = \mathbf{D}_{\phi^*}^{-1} \mathbf{B}^{*\top}$ (this follows from the symmetry of $\boldsymbol{\Omega}^*$). Ignoring this constraint allowed us to get a tractable optimization problem but caused the loss of the (asymptotic) efficiency of the estimator. This also explains why we call $\widehat{\boldsymbol{\phi}}^{\text{RML}}$ *relaxed* maximum likelihood estimator.

Proof of Proposition 2. Since $\boldsymbol{\mu}^*$ is assumed to be known and equal to zero, according to (1), we have

$$(\mathbf{S}_n \widehat{\mathbf{B}})_{jj} = \frac{1}{n} \mathbf{X}_{\bullet,j}^\top \mathbf{X} \mathbf{B}_{\bullet,j}^* = \frac{\phi_j^{*1/2}}{n} \mathbf{X}_{\bullet,j}^\top \boldsymbol{\xi}_j = \frac{\phi_j^{*1/2}}{n} (-\mathbf{X}_{\bullet,j^c} \mathbf{B}_{j^c,j}^* + \phi_j^{*1/2} \boldsymbol{\xi}_j)^\top \boldsymbol{\xi}_j.$$

Denoting $\eta_1 = -\boldsymbol{\xi}_j^\top \mathbf{X}_{\bullet, j^c} \mathbf{B}_{j^c, j}^*$, we get $(\mathbf{S}_n \widehat{\mathbf{B}})_{jj} = \frac{1}{n}(\phi_j^{*1/2} \eta_1 + \phi_j^* \|\boldsymbol{\xi}_j\|_2^2)$. Furthermore, it follows from (1) that $\mathbf{E}[(\mathbf{X}_{k, j^c} \mathbf{B}_{j^c, j}^*)^2] = \boldsymbol{\Sigma}_{jj}^* - \phi_j^*$ for each k . Since, in addition, for different k s the random variables $\mathbf{X}_{k, j^c} \mathbf{B}_{j^c, j}^*$ are independent, centered and Gaussian, we get that—in view of the independence of $\boldsymbol{\xi}_j$ and $\mathbf{X}_{\bullet, j^c}$ —the conditional distribution of η_1 given $\boldsymbol{\xi}_j$ is Gaussian with zero mean and variance $\|\boldsymbol{\xi}_j\|_2^2(\boldsymbol{\Sigma}_{jj}^* - \phi_j^*)$. Hence, the random variable $\eta = \eta_1/(\|\boldsymbol{\xi}_j\|_2(\boldsymbol{\Sigma}_{jj}^* - \phi_j^*)^{1/2})$ is standard Gaussian, independent of $\|\boldsymbol{\xi}_j\|_2^2$ and

$$(\mathbf{S}_n \widehat{\mathbf{B}})_{jj} = \frac{\sqrt{\phi_j^* \|\boldsymbol{\xi}_j\|_2^2 (\boldsymbol{\Sigma}_{jj}^* - \phi_j^*)}}{n} \eta + \frac{\phi_j^*}{n} \|\boldsymbol{\xi}_j\|_2^2.$$

This relation readily implies that $\mathbf{E}[(\mathbf{S}_n \widehat{\mathbf{B}})_{jj}] = \phi_j^*$ and

$$\mathbf{E}[(\mathbf{S}_n \widehat{\mathbf{B}})_{jj} - \phi_j^*]^2 = \mathbf{V}[(\mathbf{S}_n \widehat{\mathbf{B}})_{jj}] = \frac{\boldsymbol{\Sigma}_{jj}^* \phi_j^* + \phi_j^{*2}}{n}.$$

Furthermore, for the fourth moment, we have

$$\begin{aligned} \mathbf{E}[(\mathbf{S}_n \widehat{\mathbf{B}})_{jj} - \phi_j^*]^4 &\leq \frac{8\phi_j^{*2}(\boldsymbol{\Sigma}_{jj}^* - \phi_j^*)^2}{n^4} \mathbf{E}[\|\boldsymbol{\xi}_j\|_2^4] \mathbf{E}[\eta^4] + \frac{8\phi_j^{*4}}{n^4} \mathbf{E}[(\|\boldsymbol{\xi}_j\|_2^2 - n)^4] \\ &\leq \frac{72\phi_j^{*2}(\boldsymbol{\Sigma}_{jj}^* - \phi_j^*)^2}{n^2} + \frac{8\phi_j^{*4}}{n^4} (60n + 12n^2). \end{aligned}$$

To analyze the truncated estimator, we set $\zeta = (\mathbf{S}_n \widehat{\mathbf{B}})_{jj}$. Then $\widehat{\phi}_j^{\text{RML}} = \zeta \cdot \mathbf{1}(\zeta > 0)$ and hence,

$$\begin{aligned} \mathbf{E}[(\widehat{\phi}_j^{\text{RML}} - \phi_j^*)^2] &= \mathbf{E}[(\zeta - \phi_j^*)^2 \mathbf{1}(\zeta > 0)] + \phi_j^{*2} \mathbf{P}(\zeta \leq 0) \\ &= \mathbf{E}[(\zeta - \phi_j^*)^2] - \mathbf{E}[(\zeta - \phi_j^*)^2 \mathbf{1}(\zeta \leq 0)] + \phi_j^{*2} \mathbf{P}(\zeta \leq 0) \\ &\geq \mathbf{E}[(\zeta - \phi_j^*)^2] - \mathbf{E}[(\zeta - \phi_j^*)^4]^{1/2} \mathbf{P}(\zeta \leq 0)^{1/2}. \end{aligned}$$

We have already computed the first expectation in the right-hand side, as well as upper bounded the second one. Let us show that the probability of the event $\zeta \leq 0$ goes to zero as n increases to ∞ . This follows from the Tchebychev inequality, since $\mathbf{P}(\zeta \leq 0) = \mathbf{P}(\phi_j^* - \zeta \geq \phi_j^*) \leq \mathbf{V}[\zeta]/\phi_j^{*2} = O(1/n)$. This completes the proof of the proposition. \square

3.3. MLE taking into account the symmetry constraints

As we have seen in previous sections, the relaxed maximum likelihood estimator is suboptimal; in particular, it is less accurate than the residual variance estimator. To check that this lack of efficiency is indeed due to the relaxation of the symmetry constraints, we propose here to analyze the constrained maximum likelihood estimator in the following idealized set-up. We will consider, as in Propositions 1 and 2, that $\widehat{\mathbf{B}}$ estimates \mathbf{B}^* without error, and that⁵ there is a

⁵This assumption will be relaxed later in this subsection.

column $\mathbf{B}_{\bullet,i}^*$ in \mathbf{B}^* such that all the elements of $\mathbf{B}_{\bullet,i}^*$ are different from zero. Without loss of generality, we suppose that $i = 1$ and, consequently, for every $j \in [p]$, we have $\mathbf{B}_{j1}^* \neq 0$ which is equivalent to $\omega_{j1}^* \neq 0$. Therefore, the symmetry constraint $\mathbf{B}^* \mathbf{D}_{\phi^*}^{-1} = \mathbf{\Omega}^* = \mathbf{\Omega}^{*\top} = \mathbf{D}_{\phi^*}^{-1} \mathbf{B}^{*\top}$ implies that $\mathbf{D}_{\phi^*} \mathbf{B}^* = \mathbf{B}^{*\top} \mathbf{D}_{\phi^*}$ and, in particular, that

$$\mathbf{B}_{1j}^* \phi_1^* = \mathbf{B}_{j1}^* \phi_j^*, \quad \forall j \in [p].$$

This relation entails that in the case of known matrix \mathbf{B}^* and unknown vector ϕ^* , only the first entry of ϕ^* needs to be estimated, all the remaining entries can be computed using the first one by the formula $\hat{\phi}_j = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \hat{\phi}_1$.

Proposition 3. *Under the assumption that the rows of \mathbf{X} are i.i.d. Gaussian vectors with precision matrix $\mathbf{\Omega}^* = \mathbf{B}^* \mathbf{D}_{\phi^*}^{-1}$, the maximum likelihood estimator of ϕ^* is defined by*

$$\hat{\phi}_j^{\text{SML}} = \frac{1}{p} (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \text{trace}(\mathbf{S}_n \mathbf{B}^* \mathbf{D}_{\mathbf{B}_{\bullet,1}^*} \mathbf{D}_{\mathbf{B}_{1,\bullet}^*}^{-1}). \quad (11)$$

The quadratic risk of this estimator is given by

$$\mathbf{E}[(\hat{\phi}_j^{\text{SML}} - \phi_j^*)^2] = \frac{2}{np} \phi_j^{*2}. \quad (12)$$

Furthermore, for every $t > 0$, the following bound on the tails of the maximal error holds true:

$$\mathbf{P}\left(\max_{j \in [p]} \frac{|\hat{\phi}_j^{\text{SML}} - \phi_j^*|}{\phi_j^*} > 2\left(\frac{t + \log p}{np}\right)^{1/2} + \frac{2(t + \log p)}{np}\right) \leq 2e^{-t}.$$

Proof. To ease notation, we denote by \mathbf{D}^* the diagonal matrix whose j th element is $\mathbf{B}_{j,1}^* / \mathbf{B}_{1,j}^*$. Then, applying (8) for a given \mathbf{B}^* , the profiled Gaussian log-likelihood can be written as

$$\max_{\boldsymbol{\mu} \in \mathbb{R}^p} \mathcal{L}(\mathbf{X} | \boldsymbol{\mu}, \mathbf{B}^*, \boldsymbol{\phi}) = \frac{n}{2} \log |\mathbf{B}^*| - \frac{n}{2} \sum_{j=1}^p \{ \log(\phi_j) + (\mathbf{S}_n \mathbf{B}^*)_{jj} \phi_j^{-1} \}.$$

The goal is to maximize the right-hand side over all the vectors $\boldsymbol{\phi} \in \mathbb{R}^p$ such that $\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1}$ is a valid precision matrix.

Let us first check that under the conditions of the proposition, for $\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1}$ to be a valid precision matrix it is necessary and sufficient that $\phi_1 > 0$ and $\phi_j = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \phi_1$ for every $j \in [p]$. The necessary part follows from that fact that a precision matrix is symmetric and positive-semidefinite, which entails that $(\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1})_{1j} = (\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1})_{j1}$ and $(\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1})_{jj} = \phi_j^{-1} > 0$. Therefore, $\phi_j = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \phi_1$ and $\phi_1 > 0$. To check the sufficient part, we remark that if $\boldsymbol{\phi}$ satisfies $\phi_j = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \phi_1$ with $\phi_1 > 0$, then $\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1} = (\phi_1^* / \phi_1) \mathbf{B}^* \mathbf{D}_{\phi^*}^{-1} = (\phi_1^* / \phi_1) \mathbf{\Omega}^*$. This implies that $\mathbf{B}^* \mathbf{D}_{\boldsymbol{\phi}}^{-1}$ is symmetric and positive-semidefinite, hence a valid precision matrix.

The maximum likelihood estimator $\widehat{\phi}^{\text{SML}}$ is thus given by

$$\widehat{\phi}^{\text{SML}} \in \arg \min_{\substack{\phi \in \mathbb{R}_+^p \\ \phi_j = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \phi_1}} \sum_{j=1}^p \{ \log(\phi_j) + (\mathbf{S}_n \mathbf{B}^*)_{jj} \phi_j^{-1} \},$$

which leads to $\widehat{\phi}_1^{\text{SML}} \in \arg \min_{\phi_1 > 0} \{ p \log(\phi_1) + \phi_1^{-1} \sum_j (\mathbf{S}_n \mathbf{B}^*)_{jj} \mathbf{B}_{j1}^* / \mathbf{B}_{1j}^* \}$. The cost function of the last minimization problem is convex, since we have $(\mathbf{S}_n \mathbf{B}^*)_{jj} \mathbf{B}_{j1}^* / \mathbf{B}_{1j}^* = (\mathbf{S}_n \mathbf{B}^*)_{jj} \phi_1^* / \phi_j^* = \phi_1^* (\mathbf{S}_n \mathbf{\Omega}^*)_{jj}$. This implies that

$$\sum_j (\mathbf{S}_n \mathbf{B}^*)_{jj} \mathbf{B}_{j1}^* / \mathbf{B}_{1j}^* = \phi_1^* \text{trace}(\mathbf{S}_n \mathbf{\Omega}^*) = \phi_1^* \text{trace}(\mathbf{\Omega}^{*1/2} \mathbf{S}_n \mathbf{\Omega}^{*1/2}) \geq 0.$$

The aforementioned cost function is continuously differentiable and convex, its minimum is attained at the point where the derivative vanishes, which provides $\widehat{\phi}_1^{\text{SML}} = \frac{1}{p} \sum_j (\mathbf{S}_n \mathbf{B}^*)_{jj} \mathbf{B}_{j1}^* / \mathbf{B}_{1j}^*$. Combining with the relation $\widehat{\phi}_j^{\text{SML}} = (\mathbf{B}_{1j}^* / \mathbf{B}_{j1}^*) \widehat{\phi}_1^{\text{SML}}$, this leads to (11).

To check (12), we start by noting that

$$\widehat{\phi}_j^{\text{SML}} = \frac{1}{p} \phi_j^* \text{trace}(\mathbf{S}_n \mathbf{\Omega}^*) = \frac{1}{np} \phi_j^* \text{trace}(\mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}^{*-1}).$$

Using the well-known commutativity property of the trace operator and setting $\mathbf{Y} = \mathbf{\Sigma}^{*-1/2} \mathbf{X}^\top$, we get $\text{trace}(\mathbf{X}^\top \mathbf{X} \mathbf{\Sigma}^{*-1}) = \text{trace}(\mathbf{Y}^\top \mathbf{Y})$. Since \mathbf{X} has iid rows drawn from a $\mathcal{N}_p(0, \mathbf{\Sigma}^*)$ distribution, \mathbf{Y} has iid columns drawn from $\mathcal{N}_p(0, \mathbf{I}_p)$ distribution. Hence, the random variable $\text{trace}(\mathbf{Y}^\top \mathbf{Y}) = \sum_{j \in [p], k \in [n]} \mathbf{Y}_{jk}^2$ is distributed according to χ_{np}^2 distribution. This readily implies that $\widehat{\phi}_j^{\text{SML}}$ is an unbiased estimator of ϕ_j^* and, therefore, its quadratic risk coincides with its variance and is given by (12).

The proof of the last claim of the proposition is very similar to that of the second claim of Proposition 1. \square

Assuming that there exists $i \in [p]$ such that for any $j \in [p]$, $\omega_{ij}^* \neq 0$, put differently that the i -th node of the graph \mathcal{G}^* is connected by an edge to any other node is quite restrictive. Among other implications, it entails that the graph \mathcal{G}^* is connected which might be a strong assumption. It is therefore useful to adapt what precedes to the case where the graph \mathcal{G}^* has more than one connected component. The rest of this subsection is devoted to the description of this adaptation.

We note \mathcal{C} the set of the connected components of the graph \mathcal{G}^* . Each connected component $c \in \mathcal{C}$ is a subset of vertices of \mathcal{G}^* whose cardinality is denoted by p_c . Clearly, the sum of p_c over all $c \in \mathcal{C}$ equals p . For two vertices i and j , we will write $i \sim_{\mathcal{G}^*} j$ for indicating that they belong to the same connected component. Thus, each connected component is a class of equivalence with respect to the relation $\sim_{\mathcal{G}^*}$. Let $i \sim_{\mathcal{G}^*} j$ be two vertices from $c \in \mathcal{C}$ and let C_{ji} be a path connecting these two vertices, *i.e.*, C_{ji} is a sequence of q distinct vertices $\{v_1, \dots, v_q\}$ such that $v_1 = j$, $v_q = i$, $q \leq p_c$ and each pair (v_h, v_{h+1}) is

connected by an edge in \mathcal{G}^* . Recall that the symmetry of the precision matrix $\mathbf{\Omega}^* = \mathbf{B}^* \mathbf{D}_{\phi^*}^{-1}$ implies that $\mathbf{B}_{v_h, v_{h+1}}^* \phi_{v_h}^* = \mathbf{B}_{v_{h+1}, v_h}^* \phi_{v_{h+1}}^*$ for every $h \in [q - 1]$. This readily yields

$$\phi_j^* = \phi_i^* \prod_{1 \leq h < q} (\mathbf{B}_{v_{h+1}, v_h}^* / \mathbf{B}_{v_h, v_{h+1}}^*).$$

To ease notation, we introduce the $p \times p$ diagonal matrix $\mathbf{\Delta}_j^*$ the diagonal entries of which are defined by

$$(\mathbf{\Delta}_j^*)_{ii} = \mathbf{1}(i \sim_{\mathcal{G}^*} j) \prod_{1 \leq h < q} (\mathbf{B}_{v_{h+1}, v_h}^* / \mathbf{B}_{v_h, v_{h+1}}^*), \tag{13}$$

where $\{v_1, \dots, v_q\} = C_{ji}$ is any path connecting j to i in \mathcal{G}^* . With this notation, $\phi_j^* = (\mathbf{\Delta}_j^*)_{ii} \phi_i^*$. One can reproduce the arguments of the proof of Proposition 3 to check that the maximum likelihood estimator of ϕ^* , if \mathbf{B}^* is known (and therefore so is $\mathbf{\Delta}_j^*$), is defined by

$$\hat{\phi}_j^{\text{SML}} = \frac{1}{p_c} \text{trace}(\mathbf{\Delta}_j^* \mathbf{S}_n \mathbf{B}^*), \tag{14}$$

for j belonging to the connected component c .

Comparing the results of Propositions 1, 2 and 3, we observe that the RV-estimator outperforms the RML estimator, but—at least in the case where there is a column in \mathbf{B}^* which has only nonzero entries—they are both dominated by the maximum likelihood estimator that takes advantage of the symmetry constraints. Furthermore, using the same type of arguments as those of Proposition 3, one can check that if the vertex j of the graph \mathcal{G}^* belongs to a connected component of cardinal p_c then the risk of the MLE in the ideal case of known \mathbf{B}^* is equal to $\frac{2}{np_c} \phi_j^{*2}$. This shows that in the ideal case the MLE systematically outperforms the widely used residual variance estimator, and the gain in the risk may be huge for vertices belonging to large connected components. On the other extreme, all the three estimators discussed in the previous section coincide when the matrix \mathbf{B}^* is diagonal.

In order to apply equation (14) for estimating ϕ^* when an estimator $\hat{\mathbf{B}}$ of \mathbf{B}^* is available, we need to construct an estimator $\hat{\mathcal{G}}$ of the graph \mathcal{G}^* . We propose here an original approach for deriving $\hat{\mathcal{G}}$ from $\hat{\mathbf{B}}$. It is based on the observation that $\mathbf{B}_{ij}^* \mathbf{B}_{ji}^* = \omega_{ij}^{*2} / (\omega_{ii}^* \omega_{jj}^*)^2$, the square of the partial correlation between the i -th and j -th variables. As mentioned earlier, this quantity is always between 0 and 1 and provides a convenient rule of selection for the edges to keep in the graph. More precisely, we connect i to j if the estimated squared partial correlation $\hat{\mathbf{B}}_{ij} \hat{\mathbf{B}}_{ji}$ is larger than a prescribed threshold $t \in (0, 1)$. In our implementation, we chose (somewhat arbitrarily) the threshold $t = 0.01 \wedge n^{-1/2}$.

Note that when \mathbf{B}^* is replaced by an estimator, the right-hand side of (14) is not necessarily invariant with respect to the choice of the path connecting i to j . Therefore, even when $\hat{\mathbf{B}}$ and $\hat{\mathcal{G}}$ are fixed, if $\hat{\mathcal{G}}$ contains loops there are different ways of estimating ϕ^* based on (14) depending on how the paths are

chosen. We have tried two possible approaches: the minimum spanning tree and the shortest path tree based on the following weight function⁶ defined on the edges:

$$\mathbf{W}_{ij} = \begin{cases} \exp(-\widehat{\mathbf{B}}_{ij}\widehat{\mathbf{B}}_{ji})\mathbb{1}(\widehat{\mathbf{B}}_{ij}\widehat{\mathbf{B}}_{ji} > t), & \text{for } i \neq j, \\ 0, & \text{otherwise.} \end{cases}$$

Combining these ingredients, we get the algorithm summarized in Algorithm 1.

Algorithm 1 Estimator $\widehat{\phi}^{\text{SML}}$ based on shortest path trees or minimum spanning trees

Input: matrices \mathbf{X} and $\widehat{\mathbf{B}}$, threshold t .

Output: vector $\widehat{\phi}^{\text{SML}}$.

1: compute the matrix of weights \mathbf{W} .

2: initialize k to 1.

repeat

3: choose the node with the largest degree as root.

4: compute the shortest path tree (or the minimum spanning tree) \mathcal{T}_k from the chosen root.

5: estimate $\widehat{\phi}^{\text{SML}}$'s elements related to \mathcal{T}_k using Eq. (14)

6: remove all the nodes of the tree \mathcal{T}_k from the initial graph.

7: increment k .

until graph is empty

The rationale behind the foregoing definition of the weights and the use of the minimum spanning tree or shortest path tree algorithm is to favor the paths that are short and contain edges corresponding to large (in absolute value) partial correlations. The aim is to reduce the risk of propagating the estimation error of $\widehat{\mathbf{B}}$. We have implemented both versions of the algorithm and have observed that the version using the minimum spanning tree leads to better results. More details on the implementation and computational complexity are given in the next section.

3.4. Penalized maximum likelihood estimation

We have seen that enforcing symmetry constraints is beneficial when the matrix $\widehat{\mathbf{B}}$ has a small error, but raises intricate issues related to the graph estimation and, more importantly, path selection in the graph. A workaround to this issue is to replace the hard constraints by a penalty term that measures the degree of violation of the constraints. This provides an intermediate solution between the SML and the RML. More precisely, we propose a penalized maximum likelihood

⁶A weight equal to zero corresponds to the absence of edge.

TABLE 1

Performance of the estimators of diagonal elements of the precision matrix in Model 1. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.883 (.077)	0.399 (.036)	0.224 (.016)	1.425 (.075)	0.649 (.030)	0.374 (.022)	1.849 (.085)	0.853 (.029)	0.495 (.019)
RML	1.356 (.079)	0.786 (.040)	0.532 (.017)	2.114 (.082)	1.234 (.032)	0.841 (.022)	2.705 (.090)	1.590 (.029)	1.086 (.019)
SML	1.476 (.098)	0.805 (.040)	0.548 (.018)	2.388 (.164)	1.250 (.032)	0.852 (.021)	3.104 (.188)	1.608 (.028)	1.096 (.020)
PML	1.371 (.079)	0.792 (.041)	0.539 (.017)	2.134 (.078)	1.236 (.032)	0.846 (.021)	2.728 (.091)	1.593 (.030)	1.089 (.019)
B* estimated by square-root Lasso followed by OLS									
RV	0.726 (.079)	0.340 (.045)	0.241 (.016)	1.088 (.076)	0.616 (.051)	0.354 (.020)	1.365 (.080)	0.854 (.046)	0.443 (.018)
RML	0.726 (.079)	0.340 (.045)	0.241 (.016)	1.088 (.076)	0.616 (.051)	0.354 (.020)	1.365 (.080)	0.854 (.046)	0.443 (.018)
SML	0.807 (.082)	0.440 (.058)	0.280 (.018)	1.193 (.088)	0.793 (.066)	0.381 (.018)	1.557 (.170)	1.116 (.089)	0.468 (.018)
PML	0.737 (.074)	0.419 (.051)	0.302 (.018)	1.095 (.071)	0.722 (.052)	0.405 (.019)	1.377 (.081)	0.984 (.044)	0.494 (.019)
B* is estimated without error									
RV	0.263 (.034)	0.132 (.017)	0.081 (.012)	0.370 (.038)	0.179 (.017)	0.115 (.008)	0.455 (.038)	0.222 (.019)	0.143 (.012)
RML	0.322 (.042)	0.165 (.018)	0.104 (.013)	0.463 (.038)	0.227 (.022)	0.144 (.011)	0.562 (.040)	0.280 (.020)	0.178 (.015)
SML	0.043 (.030)	0.024 (.018)	0.010 (.010)	0.042 (.030)	0.018 (.014)	0.011 (.009)	0.042 (.037)	0.015 (.013)	0.010 (.007)
PML	0.079 (.025)	0.043 (.015)	0.023 (.007)	0.107 (.028)	0.049 (.012)	0.030 (.007)	0.128 (.027)	0.059 (.011)	0.039 (.007)

(PML) estimator of ϕ^* defined by

$$\hat{\phi}^{\text{PML}} \in \arg \min_{\phi \in (0,1]^p} \left\{ \sum_{j=1}^p \{ \log(\phi_j) + (\mathbf{S}_n \hat{\mathbf{B}})_{j,j} \phi_j^{-1} \} + \kappa \sum_{\substack{i < j \\ \hat{\mathbf{B}}_{ji} \hat{\mathbf{B}}_{ij} > t}} \frac{(\hat{\mathbf{B}}_{ji} \phi_i^{-1} - \hat{\mathbf{B}}_{ij} \phi_j^{-1})^2}{\hat{\mathbf{B}}_{ij}^2 + \hat{\mathbf{B}}_{ji}^2} \right\}, \tag{15}$$

where $\kappa > 0$ is a tuning parameter responsible for the trade-off between the likelihood and the constraint violation. The choice $\kappa = \infty$ corresponds to enforcing the symmetry constraints: its main shortcoming is that the feasible set might very well be empty. On the other extreme, when $\kappa = 0$, PML coincides with the RML. The PML estimator also coincides with the previous ones if $\hat{\mathbf{B}}$ is known to be diagonal.

Note that the parameter t appearing in the penalty term of the PML plays the same role as the one used in the SML. The definition of the feasible set in the above optimization problem is justified by the fact that we assume all the individual variances of the features to be equal to one. In other terms, the assumption $\mathbf{V}(\mathbf{X}_{1,j}) = 1$ in (1) implies that $\phi_j^* \leq 1$. Making the change of variable $\mathbf{v} = (1/\phi_j)_{j \in [p]}$, the optimization problem of Eq. (15) becomes convex with the feasible set $\mathbf{v} \in [1 + \infty)^p$ and the objective function:

$$f(\mathbf{v}) = \sum_{j=1}^p \left\{ -\log(v_j) + (\mathbf{S}_n \widehat{\mathbf{B}})_{j,j} v_j \right\} + \kappa \sum_{\substack{i < j \\ \widehat{\mathbf{B}}_{ji} \widehat{\mathbf{B}}_{ij} > t}} \frac{(\widehat{\mathbf{B}}_{ji} v_i - \widehat{\mathbf{B}}_{ij} v_j)^2}{\widehat{\mathbf{B}}_{ij}^2 + \widehat{\mathbf{B}}_{ji}^2}. \quad (16)$$

Furthermore, if we restrict the feasible set to $\mathbf{v} \in \mathcal{V} = [1, n^{1/2}]^p$, the problem becomes strongly convex. In addition, on this restricted feasible set the gradient of the objective function is Lipschitz-continuous.

It is possible to use the standard steepest gradient descent algorithm with a fixed step-size for efficiently approximating the solution $\widehat{\phi}^{\text{PML}}$. Indeed, in the optimization problem (16), if ∇f is Lipschitz-continuous with constant $L < \infty$ and strongly convex with constant $l > 0$, the gradient descent algorithm with a constant step-size $t = 2/(l + L)$ converges at a linear rate (see Nesterov (2004) for a detailed proof). Note that the convergence rate depends on L/l which is an upper bound on the condition number of the Hessian matrix $\nabla^2 f(\mathbf{v})$; this ratio should not be too high for the algorithm to converge fast. Unfortunately, the values of l and L that we manage to obtain in our problem are far too loose. That is why we resort to a steepest descent algorithm with adaptive step-size and scaled descent direction $-\nabla f(\mathbf{v}_h)/\|\nabla f(\mathbf{v}_h)\|_2$. More details on the implementation are provided in Section 4.2.

4. Experimental evaluation

In this section, we describe the experimental set-up and report the results of the numerical experiments performed on synthetic data-sets. We also provide detailed explanation of the implementation used for the symmetry-enforced and the penalized maximum-likelihood estimators. A companion R package called DESP (for estimation of Diagonal Elements of Sparse Precision-matrices) is created and uploaded on CRAN⁷.

4.1. Experiments on synthetic datasets

We conducted a comprehensive experimental evaluation of the accuracy of different estimates of diagonal elements of the precision matrix. In order to cover as many situations as possible, we used in experiments our six different forms of precision matrices along with various values for n and p . In each configuration, we considered several methods of estimating the matrix \mathbf{B}^* .

⁷<http://cran.r-project.org/web/packages/DESP/index.html>

TABLE 2

Performance of the estimators of diagonal elements of the precision matrix in Model 2. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.400 (.059)	0.125 (.020)	0.070 (.009)	0.632 (.047)	0.174 (.023)	0.094 (.011)	0.821 (.051)	0.215 (.020)	0.113 (.012)
RML	1.048 (.061)	0.508 (.020)	0.320 (.015)	1.644 (.048)	0.780 (.023)	0.491 (.014)	2.120 (.053)	0.997 (.023)	0.626 (.014)
SML	1.334 (.221)	0.539 (.028)	0.340 (.018)	2.246 (.277)	0.824 (.039)	0.520 (.023)	3.243 (.516)	1.047 (.034)	0.653 (.020)
PML	1.130 (.068)	0.530 (.020)	0.333 (.016)	1.790 (.049)	0.813 (.026)	0.508 (.016)	2.311 (.054)	1.036 (.024)	0.645 (.015)
B* estimated by square-root Lasso followed by OLS									
RV	0.247 (.053)	0.101 (.015)	0.065 (.009)	0.322 (.057)	0.129 (.019)	0.081 (.007)	0.381 (.061)	0.150 (.017)	0.095 (.009)
RML	0.247 (.053)	0.101 (.015)	0.065 (.009)	0.322 (.057)	0.129 (.019)	0.081 (.007)	0.381 (.061)	0.150 (.017)	0.095 (.009)
SML	0.329 (.107)	0.096 (.016)	0.065 (.010)	0.622 (.299)	0.129 (.021)	0.076 (.010)	0.882 (.501)	0.147 (.020)	0.090 (.011)
PML	0.247 (.068)	0.098 (.017)	0.064 (.011)	0.337 (.075)	0.125 (.021)	0.077 (.009)	0.441 (.101)	0.142 (.017)	0.089 (.011)
B* is estimated without error									
RV	0.204 (.032)	0.101 (.015)	0.065 (.008)	0.258 (.033)	0.129 (.019)	0.081 (.007)	0.300 (.030)	0.149 (.015)	0.095 (.009)
RML	0.280 (.038)	0.136 (.017)	0.086 (.011)	0.354 (.032)	0.177 (.019)	0.113 (.010)	0.429 (.038)	0.214 (.021)	0.135 (.012)
SML	0.033 (.022)	0.012 (.008)	0.008 (.007)	0.024 (.017)	0.012 (.009)	0.008 (.006)	0.027 (.019)	0.011 (.008)	0.007 (.006)
PML	0.065 (.021)	0.027 (.011)	0.019 (.006)	0.065 (.020)	0.031 (.009)	0.021 (.006)	0.073 (.022)	0.035 (.010)	0.023 (.006)

Let us first describe in a precise manner the precision matrices used in our experiments. It is worthwhile to underline here that all the precision matrices are normalized in such a way that all the diagonal entries of the corresponding covariance matrix $\Sigma^* = (\Omega^*)^{-1}$ are equal to one. To this end, we first define a $p \times p$ positive semidefinite matrix \mathbf{A} and then set $\Omega^* = (\text{diag}(\mathbf{A}^{-1}))^{\frac{1}{2}} \mathbf{A} (\text{diag}(\mathbf{A}^{-1}))^{\frac{1}{2}}$. The matrices \mathbf{A} used in the six models for which the experiments are carried out are defined as follows.

Model 1: \mathbf{A} is a Toeplitz matrix with the entries $\mathbf{A}_{ij} = 0.6^{|i-j|}$ for any $i, j \in [p]$.

Model 2: We start by defining a $p \times p$ pentadiagonal matrix with the entries

$$\bar{\mathbf{A}}_{ij} = \begin{cases} 1 & , \text{ for } |i - j| = 0, \\ -1/3 & , \text{ for } |i - j| = 1, \\ -1/10 & , \text{ for } |i - j| = 2, \\ 0 & , \text{ otherwise.} \end{cases}$$

Then, we denote by \mathbf{A} the matrix with the entries $\mathbf{A}_{ij} = (\bar{\mathbf{A}}^{-1})_{ij} \mathbb{1}(|i - j| \leq 2)$. One can check that the matrix \mathbf{A} defined in such a way is positive semidefinite.

Model 3: We set $\mathbf{A}_{ij} = 0$ for all the off-diagonal entries that are neither on the first row nor on the first column of \mathbf{A} . The diagonal entries of \mathbf{A} are

$$\mathbf{A}_{11} = p, \quad \mathbf{A}_{ii} = 2, \quad \text{for any } i \in \{2, \dots, p\},$$

whereas the off-diagonal entries located either on the first row or on the first column are $\mathbf{A}_{1i} = \mathbf{A}_{i1} = \sqrt{2}$ for $i \in \{2, \dots, p\}$.

Model 4: We introduce the integer $k = \lceil \sqrt{p} \rceil$ and define a sparse $k \times k$ matrix $\bar{\mathbf{A}}$ so that its only non-zero elements are $\bar{\mathbf{A}}_{11} = k$ and, for any $i \in [2; k]$, $\bar{\mathbf{A}}_{ii} = 2k$ and $\bar{\mathbf{A}}_{1i} = \bar{\mathbf{A}}_{i1} = \sqrt{2}$. Then, we set

$$\mathbf{A} = \begin{pmatrix} \bar{\mathbf{A}} & 0 \\ 0 & \mathbf{I}_{p-k} \end{pmatrix}.$$

Model 5: We introduce $k = \lceil \sqrt{p} \rceil$ and define a sparse $k \times k$ matrix $\bar{\mathbf{A}}$ so that its only non-zero elements are $\bar{\mathbf{A}}_{11} = 50$ and, for any $i \in [2; k]$, $\bar{\mathbf{A}}_{ii} = 5$ and $\bar{\mathbf{A}}_{1i} = \bar{\mathbf{A}}_{i1} = 5/2$. Then, similarly to previous model, we set

$$\mathbf{A} = \begin{pmatrix} \bar{\mathbf{A}} & 0 \\ 0 & \mathbf{I}_{p-k} \end{pmatrix}.$$

Model 6: We set $k = 6$, $p' = k \lceil p/k \rceil$ and define the $k \times k$ matrix $\bar{\mathbf{A}}$ as in model 5 above. Then, we build the $p' \times p'$ block-diagonal matrix \mathbf{A} by

$$\mathbf{A} = \underbrace{\begin{pmatrix} \bar{\mathbf{A}} & & 0 \\ & \ddots & \\ 0 & & \bar{\mathbf{A}} \end{pmatrix}}_{[p/k]\text{-times}}.$$

Note that, in general, the resulting precision matrix in this model is not of size $p \times p$ but of size $p' \times p'$ with $p' = 6 \lceil p/6 \rceil$. However, since in the experiments reported in this section p is always a multiple of 6, we have $p = p'$.

In this experimental evaluation, we compare the performance of the following four estimators—introduced in previous sections—of the diagonal elements of the precision matrix:

- RV corresponds to the residual variance estimator defined in Section 3.1.

- RML corresponds to the relaxed maximum likelihood estimator described by equation (10).
- SML corresponds to the symmetry-enforced maximum likelihood estimator described in Algorithm 1.
- PML corresponds to the penalized maximum likelihood estimator described by equation (15).

Note that all these algorithms need an estimator of the matrix \mathbf{B}^* to produce an estimator of the diagonal entries of the precision matrix. We conducted experiments in three different scenarios. The first scenario is when the matrix \mathbf{B}^* is estimated column-by-column by the square-root Lasso, using the penalization parameter $\lambda = \sqrt{2 \log p}$. This value for λ is commonly called the universal choice and has proved to lead to optimal theoretical results and fairly good empirical results (Dalalyan and Chen, 2012; Sun and Zhang, 2012; Dalalyan et al., 2013). The second scenario is when the matrix \mathbf{B}^* is estimated column-by-column by the ordinary least squares estimator applied to the covariates that correspond to nonzero entries of the square-root Lasso estimator⁸ with the aforementioned value of λ . Finally, the third scenario is an unrealistic one; it corresponds to the case of a known matrix \mathbf{B}^* . This scenario is included in the experimental evaluation in order to check the consistency between the theoretical and the empirical results as well as in order to better understand how the error in estimating \mathbf{B}^* impacts the quality of estimation of the diagonal entries of the precision matrix.

Thus, each configuration of our empirical study corresponds to choosing

- a model out of 6 models described above,
- a dimension $p \in \{30, 60, 90\}$,
- a sample size $n \in \{200, 800, 2000\}$,
- a method of estimating \mathbf{B}^* .

In each configuration, we computed the estimators RV, RML, SML and PML for 50 independent datasets. Using these $R = 50$ replications, we estimate the expected risk of estimating ϕ^* , $\mathbf{E}(\|\phi^* - \hat{\phi}\|_2)$, by the average $\frac{1}{R} \sum_{r=1}^R \|\phi^* - \hat{\phi}_{(r)}\|_2$. In Tables 1–6, we report these averages along with the standard deviations of the errors measured by ℓ_2 -vector norm. All the experiments were conducted in R, using the Mosek solver (see Andersen and Andersen (2000)) for computing the square-root Lasso estimator by second-order cone programming.

In the ideal case when \mathbf{B}^* is estimated without error (by itself), the empirical results reflect perfectly the theoretical results of the previous sections. The comparison of the performance of the estimators indicates that the maximum likelihood estimators SML and PML are preferable to the residual variance estimator. The maximum likelihood estimator considering symmetry constraints outperforms all the other estimators. However, in practice when $\hat{\mathbf{B}}$ is obtained by the square-root Lasso without any refinement, $\hat{\phi}^{\text{RV}}$ outperforms all the other estimators in the vast majority of configurations. Some exceptions can be observed in Models 5 and 6 (see the top part of Tables 5 and 6, where RV is slightly

⁸A discussion on the strengths and weaknesses of this estimator can be found in (Belloni and Chernozhukov, 2013; Lederer, 2014).

TABLE 3

Performance of the estimators of diagonal elements of the precision matrix in Model 3. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.273 (.042)	0.138 (.016)	0.084 (.010)	0.402 (.036)	0.194 (.014)	0.123 (.012)	0.524 (.037)	0.243 (.017)	0.150 (.011)
RML	0.509 (.062)	0.272 (.022)	0.173 (.018)	0.722 (.061)	0.395 (.026)	0.261 (.013)	0.880 (.069)	0.496 (.028)	0.321 (.014)
SML	1.080 (.132)	0.678 (.095)	0.375 (.045)	1.276 (.146)	0.802 (.075)	0.641 (.050)	1.235 (.137)	0.651 (.052)	0.454 (.029)
PML	0.509 (.062)	0.272 (.021)	0.173 (.017)	0.722 (.061)	0.395 (.026)	0.261 (.013)	0.880 (.069)	0.496 (.028)	0.322 (.014)
B* estimated by square-root Lasso followed by OLS									
RV	0.792 (.192)	0.144 (.051)	0.084 (.010)	2.251 (.203)	1.857 (.161)	0.943 (.221)	3.261 (.184)	3.815 (.157)	3.689 (.120)
RML	0.792 (.192)	0.144 (.051)	0.084 (.010)	2.251 (.203)	1.857 (.161)	0.943 (.221)	3.261 (.184)	3.815 (.157)	3.689 (.120)
SML	1.211 (.131)	0.610 (.106)	0.336 (.057)	2.515 (.189)	1.956 (.143)	1.095 (.194)	3.415 (.175)	3.832 (.152)	3.700 (.118)
PML	0.879 (.175)	0.150 (.051)	0.084 (.011)	2.366 (.207)	1.857 (.160)	0.943 (.221)	3.342 (.176)	3.816 (.157)	3.689 (.120)
B* is estimated without error									
RV	0.267 (.041)	0.138 (.016)	0.084 (.010)	0.380 (.036)	0.192 (.014)	0.122 (.011)	0.476 (.033)	0.237 (.018)	0.148 (.011)
RML	0.330 (.046)	0.163 (.016)	0.104 (.013)	0.469 (.044)	0.229 (.023)	0.151 (.013)	0.584 (.048)	0.289 (.020)	0.178 (.014)
SML	0.042 (.035)	0.019 (.013)	0.012 (.009)	0.044 (.033)	0.021 (.015)	0.011 (.007)	0.048 (.041)	0.021 (.017)	0.011 (.010)
PML	0.330 (.046)	0.163 (.016)	0.104 (.013)	0.470 (.044)	0.229 (.023)	0.151 (.012)	0.584 (.048)	0.289 (.020)	0.178 (.014)

worse than the other procedures for small sample sizes ($n = 200$). It should be, however, acknowledged that the difference of the quality between the estimators in these cases is not large enough to advocate for using RML, SML or PML. Note also that the RV estimator satisfies the following simple inequality:

$$\begin{aligned}
 (\hat{\phi}_j^{\text{RV}} - \phi_j^*)^2 &= \left(\frac{1}{n} \|\mathbf{X}\hat{\mathbf{B}}_{\bullet,j}\|_2^2 - \phi_j^* \right)^2 \\
 &\leq \frac{2}{n^2} (\|\mathbf{X}\hat{\mathbf{B}}_{\bullet,j}\|_2^2 - \|\mathbf{X}\mathbf{B}_{\bullet,j}^*\|_2^2)^2 + 2 \left(\frac{1}{n} \|\mathbf{X}\mathbf{B}_{\bullet,j}^*\|_2^2 - \phi_j^* \right)^2.
 \end{aligned}$$

The second term of the right-hand side is the error evaluated theoretically in the previous sections, while the first term can be further bounded from above

TABLE 4

Performance of the estimators of diagonal elements of the precision matrix in Model 4. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.372 (.066)	0.184 (.036)	0.113 (.023)	0.526 (.066)	0.269 (.035)	0.161 (.024)	0.655 (.076)	0.327 (.046)	0.206 (.025)
RML	0.419 (.067)	0.212 (.033)	0.134 (.020)	0.583 (.065)	0.301 (.033)	0.183 (.023)	0.722 (.074)	0.361 (.045)	0.228 (.024)
SML	0.468 (.076)	0.228 (.033)	0.144 (.020)	0.664 (.079)	0.334 (.032)	0.201 (.024)	0.843 (.095)	0.405 (.042)	0.252 (.024)
PML	0.450 (.070)	0.224 (.032)	0.142 (.020)	0.622 (.069)	0.326 (.032)	0.198 (.024)	0.763 (.073)	0.394 (.042)	0.247 (.023)
B* estimated by square-root Lasso followed by OLS									
RV	0.368 (.065)	0.182 (.036)	0.113 (.023)	0.516 (.064)	0.267 (.035)	0.160 (.024)	0.641 (.075)	0.324 (.046)	0.205 (.025)
RML	0.368 (.065)	0.182 (.036)	0.113 (.023)	0.516 (.064)	0.267 (.035)	0.160 (.024)	0.641 (.075)	0.324 (.046)	0.205 (.025)
SML	0.392 (.069)	0.191 (.037)	0.118 (.025)	0.558 (.078)	0.286 (.033)	0.173 (.025)	0.712 (.084)	0.351 (.043)	0.220 (.025)
PML	0.383 (.067)	0.188 (.037)	0.116 (.024)	0.539 (.067)	0.280 (.033)	0.169 (.024)	0.680 (.077)	0.343 (.043)	0.215 (.025)
B* is estimated without error									
RV	0.366 (.066)	0.182 (.036)	0.113 (.023)	0.515 (.065)	0.267 (.035)	0.160 (.024)	0.640 (.074)	0.324 (.046)	0.204 (.025)
RML	0.374 (.065)	0.187 (.035)	0.116 (.023)	0.524 (.066)	0.271 (.035)	0.163 (.024)	0.649 (.073)	0.330 (.046)	0.208 (.025)
SML	0.352 (.065)	0.173 (.039)	0.108 (.024)	0.500 (.066)	0.259 (.036)	0.156 (.025)	0.624 (.074)	0.316 (.046)	0.199 (.025)
PML	0.353 (.065)	0.174 (.039)	0.109 (.024)	0.500 (.066)	0.259 (.036)	0.156 (.025)	0.625 (.074)	0.317 (.046)	0.200 (.025)

by $8\left(\frac{1}{n}\{\|\mathbf{X}\widehat{\mathbf{B}}_{\bullet,j}\|_2^2 \vee \|\mathbf{X}\mathbf{B}^*_{\bullet,j}\|_2^2\}\right)\left(\frac{1}{n}\|\mathbf{X}(\widehat{\mathbf{B}}_{\bullet,j} - \mathbf{B}^*_{\bullet,j})\|_2^2\right)$. This inequality partly explains the behavior of the RV-estimator in the reported numerical results. More importantly, it shows that the error of estimating the matrix \mathbf{B}^* might have a strong impact on the quality of estimating the diagonal elements.

It is interesting to observe what happens when an additional step of estimation of \mathbf{B}^* using the ordinary least squares on the sparsity pattern provided by the square-root Lasso is performed. The impact of this step is not the same in all the models under consideration. In particular, the quality of estimation is mostly improved for all the four estimators in models 1 and 2. Furthermore, thanks to this variable selection step, the maximum-likelihood-type estimators perform nearly as well as the residual variance estimator RV. In model 3, the

TABLE 5

Performance of the estimators of diagonal elements of the precision matrix in Model 5. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.384 (.077)	0.202 (.029)	0.125 (.023)	0.543 (.060)	0.279 (.039)	0.185 (.024)	0.701 (.064)	0.342 (.040)	0.222 (.021)
RML	0.380 (.076)	0.206 (.027)	0.128 (.023)	0.539 (.060)	0.287 (.040)	0.190 (.025)	0.697 (.064)	0.352 (.041)	0.230 (.021)
SML	0.380 (.076)	0.205 (.029)	0.131 (.024)	0.539 (.060)	0.290 (.042)	0.194 (.024)	0.697 (.064)	0.353 (.041)	0.233 (.024)
PML	0.380 (.076)	0.206 (.027)	0.128 (.023)	0.539 (.060)	0.287 (.040)	0.190 (.025)	0.697 (.064)	0.352 (.041)	0.230 (.021)
B* estimated by square-root Lasso followed by OLS									
RV	0.379 (.076)	0.209 (.029)	0.130 (.025)	0.534 (.061)	0.295 (.040)	0.194 (.027)	0.693 (.064)	0.367 (.044)	0.235 (.025)
RML	0.379 (.076)	0.209 (.029)	0.130 (.025)	0.534 (.061)	0.295 (.040)	0.194 (.027)	0.693 (.064)	0.367 (.044)	0.235 (.025)
SML	0.379 (.076)	0.209 (.031)	0.134 (.026)	0.534 (.061)	0.297 (.041)	0.199 (.027)	0.693 (.064)	0.368 (.043)	0.241 (.027)
PML	0.379 (.076)	0.209 (.029)	0.130 (.025)	0.534 (.061)	0.295 (.040)	0.194 (.027)	0.693 (.063)	0.367 (.043)	0.236 (.025)
B* is estimated without error									
RV	0.384 (.075)	0.201 (.030)	0.125 (.022)	0.530 (.060)	0.275 (.038)	0.184 (.023)	0.686 (.066)	0.339 (.040)	0.221 (.022)
RML	0.383 (.076)	0.201 (.029)	0.126 (.022)	0.531 (.061)	0.277 (.037)	0.184 (.023)	0.687 (.066)	0.339 (.040)	0.221 (.022)
SML	0.347 (.078)	0.180 (.032)	0.112 (.024)	0.498 (.061)	0.257 (.042)	0.170 (.025)	0.647 (.067)	0.319 (.042)	0.206 (.023)
PML	0.383 (.076)	0.201 (.029)	0.126 (.022)	0.531 (.061)	0.277 (.037)	0.184 (.023)	0.687 (.066)	0.339 (.040)	0.221 (.022)

variable selection step deteriorates the quality of estimation in most configurations, whereas in models 4–6 this step has almost no consequence on the estimation accuracy.

The graphics of Figure 1 are drawn for Model 2 with $p = 60$. The left plot corresponds to the estimation error—measured by ℓ_2 -vector norm—as a function of the sample size in the scenario $\hat{\mathbf{B}} = \mathbf{B}^*$, whereas the central plot corresponds to the same error when \mathbf{B}^* is estimated by the OLS on the sparsity pattern furnished by the square-root Lasso. The right plot is just a zoom on the center plot. These plots illustrate the convergence to zero of the error of estimation for the estimators considered in this paper. The speed of convergence in these empirical results, as expected, is nearly $n^{-1/2}$ for fixed dimension p .

TABLE 6

Performance of the estimators of diagonal elements of the precision matrix in Model 6. The number of replications in each case is $R = 50$. More details on the experimental set-up are presented in Section 4.1.

p	30			60			90		
n	200	800	2000	200	800	2000	200	800	2000
B* estimated by square-root Lasso									
RV	0.383 (.059)	0.207 (.031)	0.140 (.018)	0.534 (.054)	0.310 (.031)	0.205 (.017)	0.651 (.057)	0.374 (.034)	0.255 (.018)
RML	0.378 (.058)	0.223 (.030)	0.157 (.020)	0.531 (.052)	0.335 (.033)	0.236 (.020)	0.648 (.055)	0.408 (.036)	0.299 (.019)
SML	0.378 (.058)	0.229 (.030)	0.169 (.022)	0.531 (.052)	0.339 (.036)	0.249 (.021)	0.649 (.055)	0.410 (.036)	0.312 (.019)
PML	0.378 (.058)	0.223 (.030)	0.157 (.020)	0.531 (.052)	0.335 (.033)	0.236 (.020)	0.648 (.055)	0.408 (.036)	0.299 (.019)
B* estimated by square-root Lasso followed by OLS									
RV	0.383 (.058)	0.245 (.030)	0.170 (.019)	0.534 (.053)	0.373 (.030)	0.262 (.024)	0.649 (.053)	0.453 (.033)	0.341 (.025)
RML	0.383 (.058)	0.245 (.030)	0.170 (.019)	0.534 (.053)	0.373 (.030)	0.262 (.024)	0.649 (.053)	0.453 (.033)	0.341 (.025)
SML	0.383 (.058)	0.251 (.027)	0.186 (.022)	0.534 (.053)	0.375 (.030)	0.281 (.023)	0.649 (.053)	0.454 (.033)	0.357 (.026)
PML	0.385 (.057)	0.245 (.029)	0.170 (.019)	0.534 (.053)	0.373 (.030)	0.262 (.024)	0.650 (.053)	0.453 (.033)	0.341 (.025)
B* is estimated without error									
RV	0.408 (.068)	0.210 (.030)	0.141 (.018)	0.569 (.068)	0.309 (.031)	0.205 (.018)	0.697 (.063)	0.370 (.030)	0.251 (.018)
RML	0.411 (.070)	0.212 (.030)	0.142 (.019)	0.578 (.067)	0.313 (.033)	0.208 (.018)	0.702 (.064)	0.372 (.031)	0.254 (.018)
SML	0.182 (.057)	0.097 (.023)	0.061 (.020)	0.277 (.064)	0.142 (.033)	0.094 (.022)	0.311 (.073)	0.178 (.030)	0.110 (.019)
PML	0.411 (.070)	0.212 (.030)	0.142 (.019)	0.578 (.067)	0.313 (.033)	0.208 (.018)	0.702 (.064)	0.372 (.031)	0.254 (.018)

4.2. Details on the implementation

Symmetry-enforced maximum likelihood As we explained earlier, the product structure of the term Δ_j^* in (13) may cause the amplification of the estimation error when passing from $\hat{\mathbf{B}}$ to $\hat{\phi}$. In order to reduce as much as possible this phenomenon, we suggested to choose the path \mathcal{C} by minimizing its length. In addition, the fact that some entries of \mathbf{B}^* appear in the denominator of Δ_j^* , make it unsuitable to include in \mathcal{C} edges corresponding to small values of $\hat{\mathbf{B}}_{ij}$. The combination of these two arguments suggests to define edge weights as decreasing functions of $\hat{\mathbf{B}}_{ij}$ and to look for paths that somehow minimize the overall weight defined as the sum of the weights of the edges contained in \mathcal{C} .

The two versions of the SML algorithm that have been implemented and tested in this work make use of the minimum spanning tree (MST) and the shortest path tree in the step of determining the way of computation the elements of $\hat{\phi}$ belonging to a connected component \mathcal{C} of the graph $\hat{\mathcal{G}}$. A MST of \mathcal{C} is a tree that spans \mathcal{C} and has the smallest total weight among all the spanning trees of \mathcal{C} . The shortest path tree having a given node r as a root is a spanning tree \mathcal{T} of \mathcal{C} such that for any node $j \in \mathcal{C}$ the weight of the path from j to r in \mathcal{T} is the smallest among the weights of all possible paths from j to r in \mathcal{C} .

We have used the Kruskal (Kruskal, 1956) algorithm for finding the MST and the Jarnik-Prim-Dijkstra algorithm (Jarník, 1930; Prim, 1957; Dijkstra, 1959) for the shortest path tree. The worst-case computational complexities of the construction of these trees are the following (Cormen et al., 2009). When the graph \mathcal{G} has p nodes and q edges, the Kruskal algorithm runs in $O(q \log p)$ time. Its output is a set of MSTs per connected component. The version of the SML based on the shortest path tree requires $O(p + q)$ operations to find the connected components. In a connected component having p_c nodes and q_c edges, the node of largest degree can be obtained in $O(q_c)$ operations, while the computational complexity of finding the shortest paths from a node to all the others is $O(q_c \log(p_c))$. Therefore, determining a shortest path tree per connected component has a complexity of $O(p + q \log(p))$, or $O(sp \log(p))$ where s is the maximal degree of a node of $\hat{\mathcal{G}}$. Thus, the computational complexities of the two versions of the SML estimator are comparable and, at most, of the order $O(sp \log(p))$.

In our experiments, we have also tried⁹ a third version consisting in computing the shortest path trees from every node of a connected component and then choosing the one with the minimal overall weight, rather than first choosing the root as the node having largest degree. Several other variants have been tested as well, but the simplest version based on choosing the MST has lead to the best empirical results.

Penalized maximum likelihood As mentioned earlier, the PML estimator is computed by solving the optimization problem (16). We implement a steepest descent algorithm with adaptive step-size and scaled descent direction $-\nabla f(\mathbf{v}_h)/\|\nabla f(\mathbf{v}_h)\|_2$. At each iteration, one common adaptation for every coordinate of the descent direction is performed. If the objective function increases, the current iteration is done again with a halved step-size. On the opposite, if the objective function decreases, the step-size is increased by a constant factor for the next iteration.

Mathematically speaking, the update operations for our gradient descent algorithm are

$$\mathbf{v}_0 = \mathbf{1}, \quad \mathbf{v}_{h+1} = \mathbf{v}_h + t_h \mathbf{u}_h, \quad h = 0, 1, 2, \dots, \quad (17)$$

where the descent direction is $\mathbf{u}_h = -\nabla f(\mathbf{v}_h)/\|\nabla f(\mathbf{v}_h)\|_2$ and t_h is the step-size. Thanks to the convexity, the convergence of this algorithm is guaranteed

⁹We used the package RBGL of R (Long et al., 2016) for various algorithms related to weighted graphs.

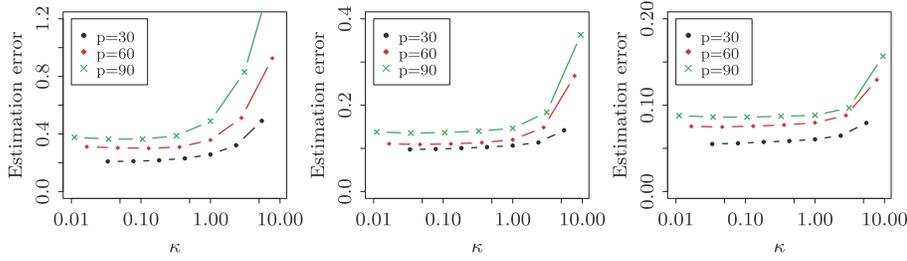


FIG 2. The estimation error of the PML as a function of κ . The plots are obtained for the synthetic experiment of Model 2 with various values of p and for $n = 200$ (left), $n = 800$ (middle) and $n = 2000$. Please note that the limits of the y-axis are not the same in the three plots and that the x-axis is presented in logarithmic scale.

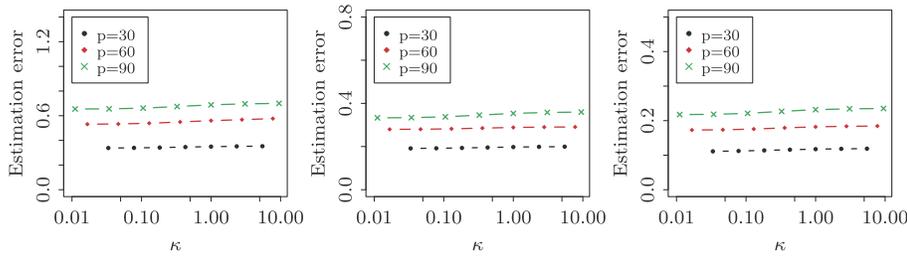


FIG 3. The estimation error of the PML as a function of κ . The plots are obtained for the synthetic experiment of Model 4 with various values of p and for $n = 200$ (left), $n = 800$ (middle) and $n = 2000$. Please note that the limits of the y-axis are not the same in the three plots and that the x-axis is presented in logarithmic scale.

for any starting point \mathbf{v}_0 . The step-size is updated at each iteration according to the following rule:

$$t_{h+1} = \begin{cases} 1.2 \times t_h, & \text{for } f(\mathbf{v}_{h+1}) < f(\mathbf{v}_h), \\ 0.5 \times t_h, & \text{otherwise.} \end{cases}$$

The multiplicative factors we use for adaptive step-size are those propose by Riedmiller and Braun (1992) for the Rprop algorithm. We stop iterating when the gradient magnitude measured in the ℓ_2 -norm is below a certain level (10^{-5} in our experiments) or when the limit of 5000 iterations is attained.

For the choice of the tuning parameter κ , we did a cross-validation by choosing a geometric grid over the values of κ ranging from $1/p$ to \sqrt{p} . The results, for Models 2 and 4, are plotted in Fig. 2 and 3, respectively. We can clearly see that there is a large interval of values of κ for which the error is nearly minimal. Based on this observation, we chose $\kappa = \frac{1}{3}\sqrt{\log p}$ for all the numerical experiments reported in Tables 1–6.

5. Conclusion

This paper introduces three estimators of the diagonal entries of a sparse precision matrix when n iid copies of a Gaussian vector with this precision matrix are observed. The properties of these estimators are discussed and compared with those of the commonly used residual variance estimator. At a theoretical level, an interesting finding is that the naive maximum likelihood estimator (MLE) that does not take into account the symmetry constraints has a significantly larger risk than the residual variance estimator and, hence, is not optimal even asymptotically. The symmetry-enforced MLE and the penalized MLE circumvent this drawback and are shown in all numerical experiments to outperform the residual variance estimator when the matrix \mathbf{B}^* is known. Similar but unreported results are obtained when the estimators of the diagonal entries use a noisy matrix $\hat{\mathbf{B}} = \mathbf{B}^* + \mathbf{\Xi}$, provided the noise matrix $\mathbf{\Xi}$ has iid Gaussian entries with zero mean and small variance. However, in a more realistic situation when \mathbf{B}^* is estimated by the square-root Lasso or by the ordinary least squares conducted over the submodel selected by the square-root Lasso, the accuracies of the four estimators of the diagonal entries become comparable with a slight advantage for the residual variance estimator.

We would like also to mention the introduction of a novel and simple method of estimating partial correlations and of symmetrizing the precision matrix estimator derived from the nonsymmetric matrix $\hat{\mathbf{B}}$. It is based on the observation that the square of the partial correlation between i -th and j -th variables is equal to $\mathbf{B}_{ij}^* \mathbf{B}_{ji}^*$.

In the future, it would be interesting to look for an estimator of \mathbf{B}^* which is more accurate than the square-root Lasso and could hopefully—in combination with the symmetry-enforced MLE or the penalized MLE—lead to better precision matrix estimate than the one obtained by the association of the square-root Lasso and the residual variance estimator. Another appealing avenue for future research is the investigation of the case when the matrix \mathbf{X} is observed with an error. Recent papers (Rosenbaum and Tsybakov, 2013; Belloni et al., 2014b) may provide valuable guidance for accomplishing this task.

Acknowledgments

The work of the second author was partially supported by the grant Investissements d’Avenir (ANR-11-IDEX-0003/Labex Ecodec/ANR-11-LABX-0047) and the chair “LCL/GENES/Fondation du risque, Nouveaux enjeux pour nouvelles données”.

References

- E. D. Andersen and K. D. Andersen. The mosek interior point optimizer for linear programming: An implementation of the homogeneous algorithm. In *High Performance Optimization*, pages 197–232. 2000. [MR1748773](#)

- T. W. Anderson. *An introduction to multivariate statistical analysis*. Wiley Series in Probability and Statistics. Wiley-Interscience [John Wiley & Sons], Hoboken, NJ, third edition, 2003. [MR1990662](#)
- O. Banerjee, L. El Ghaoui, and A. d'Aspremont. Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data. *J. Mach. Learn. Res.*, 9:485–516, June 2008. [MR2417243](#)
- A. Belloni and V. Chernozhukov. Least squares after model selection in high-dimensional sparse models. *Bernoulli*, 19(2):521–547, May 2013. [MR3037163](#)
- A. Belloni, V. Chernozhukov, and L. Wang. Square-root lasso: Pivotal recovery of sparse signals via conic programming. *Biometrika*, 98(4):791–806, 2011. [MR2860324](#)
- A. Belloni, V. Chernozhukov, and L. Wang. Pivotal estimation via square-root Lasso in nonparametric regression. *Ann. Statist.*, 42(2):757–788, 2014a. [MR3210986](#)
- A. Belloni, M. Rosenbaum, and A. B. Tsybakov. An l_1, l_2, l_∞ -regularization approach to high-dimensional errors-in-variables models. Technical Report CREST, arxiv:1412.7216, 2014b.
- P. Bühlmann and S. A. van de Geer. *Statistics for High-dimensional data: methods, theory and applications*. Springer series in statistics. Springer-Verlag Berlin Heidelberg, 2011. [MR2807761](#)
- T. Cai, W. Liu, and X. Luo. A constrained L1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, 106:594–607, February 2011. [MR2847973](#)
- T. Cai and W. Liu. A direct estimation approach to sparse linear discriminant analysis. *J. Amer. Statist. Assoc.*, 106(496):1566–1577, 2011. [MR2896857](#)
- T. Cai, W. Liu, and H. Zhou. Estimating sparse precision matrix: Optimal rates of convergence and adaptive estimation. *Ann. Statist.*, 44(2):455–488, 2016. [MR3476606](#)
- E. Candes and T. Tao. The dantzig selector: Statistical estimation when p is much larger than n . *Ann. Statist.*, 35(6):2313–2351, December 2007. [MR2382644](#)
- T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to algorithms*. MIT Press, Cambridge, MA, third edition, 2009. [MR2572804](#)
- A. S. Dalalyan and Y. Chen. Fused sparsity and robust estimation for linear models with unknown variance. In *Advances in Neural Information Processing Systems 25: NIPS*, pages 1268–1276, 2012.
- A. S. Dalalyan, M. Hebiri, K. Meziani, and J. Salmon. Learning heteroscedastic models by convex programming under group sparsity. In *Journal of Machine Learning Research – W & CP 28(3) (ICML 2013)*, page 379–387, 2013.
- E. W. Dijkstra. A note on two problems in connexion with graphs. *Numer. Math.*, 1:269–271, 1959. [MR0107609](#)
- R. A. Fisher. The use of multiple measurements in taxonomic problems. *Annals of Eugenics*, 7(2):179–188, 1936.
- J. Friedman, T. Hastie, and R. Tibshirani. Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, 9(3):432–441, July 2008.
- I. A. Ibragimov and R. Z. Has'minskiĭ. *Statistical estimation*, volume 16 of *Appli-*

- cations of Mathematics*. Springer-Verlag, New York-Berlin, 1981. Asymptotic theory, Translated from the Russian by Samuel Kotz. [MR0620321](#)
- V. Jarník. *O jistém problému minimálním: (Z dopisu panu O. Borůskovi)*. Práce Moravské přírodovědecké společnosti. Mor. přírodovědecká společnost, 1930.
- J. B. Kruskal, Jr. On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Amer. Math. Soc.*, 7:48–50, 1956. [MR0078686](#)
- J. Lafferty, H. Liu, and L. Wasserman. Sparse nonparametric graphical models. *Statist. Sci.*, 27(4):519–537, 2012. [MR3025132](#)
- B. Laurent and P. Massart. Adaptive estimation of a quadratic functional by model selection. *Ann. Statist.*, 28(5):1302–1338, 2000. [MR1805785](#)
- S. L. Lauritzen. *Graphical models*, volume 17 of *Oxford Statistical Science Series*. The Clarendon Press, Oxford University Press, New York, 1996. Oxford Science Publications. [MR1419991](#)
- L. Le Cam and G. L. Yang. *Asymptotics in statistics*. Springer Series in Statistics. Springer-Verlag, New York, second edition, 2000. Some basic concepts. [MR1784901](#)
- J. Lederer. Trust, but verify: Benefits and pitfalls of least-squares refitting in high dimensions. Technical report, arXiv:1306.0113, 2014.
- H. Liu and L. Wang. Tiger: A tuning-insensitive approach for optimally estimating large undirected graphs. Technical report, arxiv:1412.7216, 2012.
- H. Liu, J. Lafferty, and L. Wasserman. The nonparanormal: Semiparametric estimation of high dimensional undirected graphs. *J. Mach. Learn. Res.*, 10: 2295–2328, 2009. [MR2563983](#)
- L. Long, V. Carey, and R. Gentleman. *RBGL: An interface to the BOOST graph library*, 2016. URL <http://www.bioconductor.org>.
- G. Marsaglia. Conditional means and covariances of normal variables with singular covariance matrix. *Journal of the American Statistical Association*, 59 (308):1203–1204, 1964.
- N. Meinshausen and P. Bühlmann. High-dimensional graphs and variable selection with the lasso. *Ann. Statist.*, 34(3):1436–1462, June 2006. [MR2278363](#)
- Y. Nesterov. *Introductory lectures on convex optimization: A basic course*. Applied optimization. Kluwer Academic Publ., Boston, Dordrecht, London, 2004. ISBN 9781402075537. [MR2142598](#)
- R. C. Prim. Shortest connection networks and some generalizations. *Bell System Technology Journal*, 36:1389–1401, 1957.
- M. Riedmiller and H. Braun. Rprop – a fast adaptive learning algorithm. Technical report, Proc. of ISCIS VII), Universitat, 1992.
- M. Rosenbaum and A. B. Tsybakov. *Improved matrix uncertainty selector*, volume Volume 9 of *Collections*, pages 276–290. Institute of Mathematical Statistics, 2013. [MR3202640](#)
- N. Städler, P. Bühlmann, and S. van de Geer. ℓ_1 -penalization for mixture regression models. *TEST*, 19(2):209–256, 2010a. [MR2677722](#)
- N. Städler, P. Bühlmann, and S. van de Geer. Rejoinder: ℓ_1 -penalization for mixture regression models. *TEST*, 19(2):280–285, 2010b. [MR2677728](#)

- T. Sun and C-H. Zhang. Scaled sparse linear regression. *Biometrika*, 99(4):879–898, September 2012. [MR2999166](#)
- T. Sun and C-H. Zhang. Sparse matrix inversion with scaled lasso. *J. Mach. Learn. Res.*, 14:3385–3418, November 2013. [MR3144466](#)
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society (Series B)*, 58:267–288, 1996. [MR1379242](#)
- M. Yuan. High dimensional inverse covariance matrix estimation via linear programming. *J. Mach. Learn. Res.*, 11:2261–2286, January 2010. [MR2719856](#)
- M. Yuan and Y. Lin. Model selection and estimation in the gaussian graphical model. *Biometrika*, 94(1):19–35, 2007. [MR2367824](#)