

Estimating the correlation in network disturbance models

A. D. BARBOUR

Institut für Mathematik, Universität Zürich, Winterthurerstrasse 190, CH-8057 Zürich, Switzerland

AND

GESINE REINERT[†]

Department of Statistics, University of Oxford, 24–29 St Giles', Oxford OX1 3LB, UK

[†]Corresponding author. Email: reinert@stats.ox.ac.uk

Edited by: James Gleeson

[Received on 11 December 2020; editorial decision on 18 August 2021; accepted on 24 August 2021]

The network disturbance model of P. Doreian (1989), expresses the dependency between observations taken at the vertices of a network by modelling the correlation between neighbouring vertices, using a single correlation parameter ρ . It has been observed that estimation of ρ in *dense* graphs, using the method of maximum likelihood, leads to results that can be both biased and very unstable. In this article, we sketch why this is the case, showing that the variability cannot be avoided, no matter how large the network. We also propose a more intuitive estimator of ρ , which shows little bias. The related network effects model is briefly discussed.

Keywords: network disturbance; network autocorrelation; maximum likelihood.

1. Introduction

The classical Gaussian linear regression model can be written as

$$y = X\beta + \varepsilon. \quad (1.1)$$

Here, the n -dimensional vector y of data is described as the sum of a structural element $X\beta$, where X is a known $n \times m$ matrix of full rank m , $\beta = (\beta_1, \dots, \beta_m)^\top$ is an m -vector of unknown parameters, to be estimated, and the elements $\varepsilon_1, \dots, \varepsilon_n$ of the n -vector ε are realizations of independent normally distributed random variables $\varepsilon_1, \dots, \varepsilon_n$, each with zero mean and variance σ^2 . Here and throughout, we conventionally use bold face to denote random elements, and italic for their realizations.

If the indices $i = 1, 2, \dots, n$ represent the vertices of a simple undirected graph, whose (symmetric) adjacency matrix is denoted by A , it may be reasonable to suppose that the errors ε_i and ε_j at neighbouring vertices i and j are correlated, instead of being independent. A similar idea lies at the basis of time series analysis, where the error at a given time point is assumed to depend on the errors at the closest previous time points. In analogy with the moving average models of time series analysis, the *Network Disturbance Model* of [1] (see also [2, (4.6)]) supposes that $\mathbf{y} = X\beta + \boldsymbol{\varepsilon}$, where the errors $\boldsymbol{\varepsilon}$ are now implicitly given by the equation

$$\varepsilon_i = \rho \sum_{j=1}^n W_{ij} \varepsilon_j + v_i, \quad 1 \leq i \leq n, \quad (1.2)$$

or more succinctly, by

$$K_\rho \boldsymbol{\varepsilon} = \boldsymbol{v}, \quad \text{where} \quad K_\rho := I - \rho W. \quad (1.3)$$

Here, the weight matrix W is known, the parameter ρ is unknown, and the random variables $\boldsymbol{v}_1, \dots, \boldsymbol{v}_n$ are independent normal random variables with mean zero and variance σ^2 . The weight $W_{ij} \geq 0$ can be thought of as a (relative) measure of the influence of vertex j on vertex i and is assumed to take the value zero if i and j are not neighbours, that is, if $A_{ij} = 0$; we also assume that $W_{ii} = 0$ for each i . The parameter ρ can be varied, and $|\rho|$ mirrors the overall strength of the dependence between the observations at neighbouring vertices.

The network disturbance model can be estimated by using the method of maximum likelihood (MLE); details of the procedure are given in Section 2. In classical settings, including repeated independent observations from a common underlying distribution, the maximum likelihood estimator is asymptotically unbiased, as the number of observations tends to infinity. Here, the observations are dependent, and the standard asymptotic conclusions need no longer be valid. Indeed, it has been widely observed in simulation studies that, as the underlying network becomes denser, and the corresponding weight matrix W becomes less sparse, the maximum likelihood estimator $\hat{\rho}$ of ρ becomes more variable and more negatively biased. This has been extensively documented, largely in the context of the related network effects models of [1], defined in (1.4) below, for instance in [3–7]. Methods for correcting the bias of the MLE, using asymptotics appropriate in settings where consistent estimation is possible, have since been proposed in [8–10]. Other methods of estimation have been developed by [11], who introduce a generalized method of moments estimator, by [12], with a Bayesian approach, and by [13], whose techniques are based on Hodges–Lehmann estimators. In [7], theoretical investigation shows that, at least for networks very close to the complete graph, the MLE may lead to misleading results: see the final example in Section 4.4 for a fuller explanation, that makes clear why no estimator of ρ can work well in such circumstances.

In this article, we examine the problem in greater depth, showing why the MLE can be expected to be biased and lacking in precision, whenever the underlying network is dense. In Section 3, we use Cramér–Rao theory to show that, if the underlying network is sufficiently dense, then *any* (more or less unbiased) estimator of ρ must necessarily be unreliable. The bound is expressed as a lower bound (3.7) on the variance of any (unbiased) estimator, and is given explicitly as a function of the matrix W of weights and of the true value of ρ . The structure of the network enters by way of W . For instance, the choice $W_{ij} := A_{ij} / \deg(i)$ makes explicit the dependence of the weights on the adjacency matrix A of the network. We then, in Section 4, examine the behaviour of the maximum likelihood estimator in enough detail to explain the empirical observations. In particular, we are able to show that the MLE of ρ is biased when the network is sufficiently dense. In ordinary least squares estimation, the estimator of β has small variance if the smallest eigenvalue of the matrix $X^\top X$ is large. Here, it is possible, for rather special graphs, that the MLE of β may still have substantial variance, even when this condition is satisfied; see Examples 1 and 2 in Section 4.4.

The exact formulae for the bias and variance of the MLE are in general indigestible, but we discuss some special cases in which it is possible to estimate their magnitudes. In Section 5, we consider a simpler and more intuitive estimator of ρ , that has better properties as regards bias, but is still lacking in precision in dense graphs, as it must be, in view of the lower bound given in Section 3. Its properties are illustrated by some simulations. An R function for computing the estimator, as well an empirical measure of its precision, is available on request from the authors.

There is a related model, the *network effects model* of [1] (see also [2, (4.1)], which takes the form

$$K_\rho y = X\beta + \nu, \quad (1.4)$$

where the components of ν are realizations of independent normal random variables with mean zero and variance σ^2 , and K_ρ is as defined in (1.3). This differs from the network disturbance model because, rewriting it as

$$y = K_\rho^{-1}X\beta + \varepsilon, \quad (1.5)$$

the errors in y have the same dependence as before, but the *model structure* explaining y also changes with ρ . The limitations as regards estimating ρ in dense networks are analogous, and for similar reasons. We discuss the model briefly in Section 6.

The main conclusions of the article are sketched in Section 7, and additional calculations can be found in the Appendix.

2. Details of the model

The maximum likelihood estimates of the parameters in the simple linear regression model $\mathbf{y} = X\beta + \boldsymbol{\varepsilon}$ of (1.1), with independent and identically distributed errors, are given by

$$\hat{\beta} = \hat{\beta}(y) = (X^\top X)^{-1}X^\top y; \quad \hat{\sigma}^2 = \hat{\sigma}^2(y) = n^{-1}y^\top (I - H(X))y, \quad (2.1)$$

where

$$H(X) := X(X^\top X)^{-1}X^\top \quad (2.2)$$

is symmetric and idempotent, and represents the orthogonal projection from \mathbb{R}^n onto the linear span of the columns of X . Writing

$$M_0: \mathbf{y} = X\beta + \boldsymbol{\varepsilon} \quad (2.3)$$

to denote the probabilistic model from which the data are sampled, the random quantities $\hat{\beta}(\mathbf{y})$ and $\hat{\sigma}^2(\mathbf{y})$ have distributions given by

$$\hat{\beta}(\mathbf{y}) \sim \text{MVN}_m(\beta, \sigma^2(X^\top X)^{-1}); \quad \hat{\sigma}^2(\mathbf{y}) \sim n^{-1}\sigma^2\chi_{n-m}^2, \quad (2.4)$$

where $\text{MVN}_m(\mu, \Sigma)$ denotes the m -dimensional multivariate normal distribution with mean μ and covariance matrix Σ , and χ_l^2 denotes the chi-squared distribution with l degrees of freedom. The two estimators are also independent of one another (see [14, Section 4.b]).

In the network disturbance model $\mathbf{y} = X\beta + \boldsymbol{\varepsilon}$, the errors $\boldsymbol{\varepsilon}$ are given by (1.3), with dependence deriving from the weight matrix W . It is often convenient to assume that the row sums of W are all equal to 1, in which case the quantity $\sum_{j=1}^n W_{ij}\boldsymbol{\varepsilon}_j$ represents an average of the errors at the neighbours of vertex i ; then ρ would typically be smaller than 1 in modulus, so as to ensure that the average influence of the neighbours is not larger than the size of a typical error. A simple choice for W_{ij} is $A_{ij}/\deg(i)$, where $\deg(i)$

denotes the degree of the vertex i . If there are isolated vertices i , the corresponding row sums cannot be made equal to 1, since then $W_{ij} = 0$ for all $j \neq i$.

The implicit representation (1.3) of the errors can be written as a direct formula for $\boldsymbol{\varepsilon}$ in terms of \mathbf{v} , as

$$\boldsymbol{\varepsilon} = K_\rho^{-1} \mathbf{v}, \quad (2.5)$$

provided that the matrix K_ρ can be inverted. The expansion

$$K_\rho^{-1} = (I - \rho W)^{-1} = I + \sum_{l \geq 1} \rho^l W^l \quad (2.6)$$

makes sense provided that the sum $\sum_{l \geq 1} \rho^l W^l x$ converges absolutely for each $x \in \mathbb{R}^n$; that is, if $|\rho|$ is less than $1/r(W)$, where $r(W)$ is the spectral radius of W . In particular, if W has all of its row sums equal to 1, the vector $\mathbf{1}$ with all components equal to 1 is a right eigenvector with eigenvalue 1, implying that $r(W) = 1$ (see [15, Corollary 8.1.30, p. 493]). Then

$$\boldsymbol{\varepsilon} = K_\rho^{-1} \mathbf{v} = \mathbf{v} + \sum_{l \geq 1} \rho^l W^l \mathbf{v}, \quad |\rho| < 1,$$

represents $\boldsymbol{\varepsilon}_i$ as a weighted sum of the contributions of the different elements \mathbf{v}_j , averaged along paths in the network from j to i , and with overall weight ρ^l for paths of length l (possibly negative, if $\rho < 0$). If $|\rho| > 1/r(W)$, this interpretation no longer holds, since the expansion (2.6) is no longer well defined. The resulting model is then not so easy to understand, since, at least locally, the cumulative effect on the observation at i from those at distance d from i increases geometrically with d . In time series analysis, the analogue is that of processes that are ‘explosive’, and the analogue of the condition $|\rho| < 1/r(W)$ is used to ensure that a time series is stationary. For this reason, we shall assume in what follows that $|\rho| < 1/r(W)$, so that the inverse of K_ρ exists, and is explicitly given by (2.6).

The model $\mathbf{y} = X\boldsymbol{\beta} + \boldsymbol{\varepsilon}$ with errors given by (1.2) is simple to analyse if the value of ρ is known, since it can be written in the form

$$M_\rho: \quad K_\rho \mathbf{y} = K_\rho X \boldsymbol{\beta} + \mathbf{v}, \quad (2.7)$$

with the components of \mathbf{v} being independent normal $\mathcal{N}(0, \sigma^2)$ random variables, and can thus be fitted as a standard linear regression based on data $\mathbf{y}' := K_\rho \mathbf{y}$. Equation (2.1) then gives the estimates

$$\begin{aligned} \hat{\boldsymbol{\beta}}_\rho &= \hat{\boldsymbol{\beta}}_\rho(\mathbf{y}) = ((K_\rho X)^\top K_\rho X)^{-1} (K_\rho X)^\top K_\rho \mathbf{y}; \\ \hat{\sigma}_\rho^2 &= \hat{\sigma}_\rho^2(\mathbf{y}) = n^{-1} \{K_\rho(\mathbf{y} - X \hat{\boldsymbol{\beta}}_\rho)\}^\top \{K_\rho(\mathbf{y} - X \hat{\boldsymbol{\beta}}_\rho)\} = n^{-1} (K_\rho \mathbf{y})^\top (I - H_\rho) K_\rho \mathbf{y}, \end{aligned} \quad (2.8)$$

where, with H as in (2.2),

$$H_\rho := H(K_\rho X) \quad (2.9)$$

and so $K_\rho X \hat{\boldsymbol{\beta}}_\rho = H_\rho K_\rho \mathbf{y}$. The residuals after fitting M_ρ are thus

$$\hat{\boldsymbol{\varepsilon}}_\rho(\mathbf{y}) := \mathbf{y} - X \hat{\boldsymbol{\beta}}_\rho(\mathbf{y}) = K_\rho^{-1} (I - H_\rho) K_\rho \mathbf{y}. \quad (2.10)$$

The distributions of $\hat{\beta}_\rho(\mathbf{y})$ and $\hat{\sigma}_\rho^2(\mathbf{y})$ under the model M_ρ are given by

$$\hat{\beta}_\rho(\mathbf{y}) \sim \text{MVN}_m(\beta, \sigma^2((K_\rho X)^\top K_\rho X)^{-1}); \quad \hat{\sigma}_\rho^2(\mathbf{y}) \sim n^{-1}\sigma^2\chi_{n-m}^2, \quad (2.11)$$

and the two estimators are independent of one another. In particular, both $\hat{\beta}$ and $\hat{\sigma}^2$ have precision much as in simple linear regression, and (under reasonable assumptions about X) can be expected to become more accurate as n increases. If the true model is $M_{\rho'}$, for $\rho' \neq \rho$, the expectation of $\hat{\beta}_\rho(\mathbf{y})$ is still equal to β , and, although the variances of its components may be larger, they are typically comparable to those on the true model, again under reasonable conditions; see Appendix A.4 for more detail.

If ρ is not known, the estimates (2.8) still maximize the log-likelihood

$$\ell(\theta; \mathbf{y}) := -\frac{n}{2}\{\log(\sigma^2) + \log(2\pi)\} - \frac{1}{2\sigma^2}\{K_\rho(\mathbf{y} - X\beta)\}^\top \{K_\rho(\mathbf{y} - X\beta)\} + \log \det K_\rho \quad (2.12)$$

of the model M_ρ of (2.7), for each given value of ρ , where $\theta := (\beta, \sigma^2, \rho)$ denotes the full set of parameters; see [1, Equation (1.3)]. The resulting maximum of the log-likelihood, for fixed ρ , is given by

$$-\frac{n}{2}\{\log(\hat{\sigma}_\rho^2(\mathbf{y})) + \log(2\pi) + 1\} + \log \det K_\rho; \quad (2.13)$$

this can be deduced by observing that, with β replaced by $\hat{\beta}_\rho$, the middle term of (2.12) reduces to $-n/2$, in view of (2.8). Indeed, the first term in (2.13) is the maximum of the likelihood for ordinary least squares applied to the data $\mathbf{y}' := K_\rho \mathbf{y}$, and the determinant enters because the probability densities $f_{\mathbf{y}'}$ of \mathbf{y}' and $f_{\mathbf{y}}$ of \mathbf{y} are related by $f_{\mathbf{y}'}(\mathbf{v}) = \det K_\rho f_{\mathbf{y}}(K_\rho \mathbf{v})$. Then, ρ can be chosen to (numerically) maximize the expression in (2.13), or, equivalently, to minimize

$$F(\rho; \mathbf{y}) := \log(\hat{\sigma}_\rho^2(\mathbf{y})) - \frac{n}{2} \log \det K_\rho, \quad |\rho| < 1/r(W). \quad (2.14)$$

3. A general limit to precision

The ρ -component of the score function is obtained by differentiating the log-likelihood (2.12) with respect to ρ . Note that $\frac{d}{d\rho} K_\rho = -W$, and that, if the eigenvalues of W are denoted by $\lambda_1, \dots, \lambda_n$, we have $\det K_\rho = \prod_{i=1}^n (1 - \rho\lambda_i)$. This gives the ρ -component of the score function as

$$\begin{aligned} s^{(\rho)}(\theta; \mathbf{y}) &:= \frac{\partial}{\partial \rho} \ell(\theta; \mathbf{y}) = \frac{1}{\sigma^2} (\mathbf{y} - X\beta)^\top K_\rho^\top W (\mathbf{y} - X\beta) - \sum_{l=1}^n \frac{\lambda_l}{1 - \rho\lambda_l} \\ &= \frac{1}{\sigma^2} (\{K_\rho(\mathbf{y} - X\beta)\}^\top Z_\rho \{K_\rho(\mathbf{y} - X\beta)\}) - \text{Tr}\{Z_\rho\}, \end{aligned} \quad (3.1)$$

where we define the matrix

$$Z_\rho := WK_\rho^{-1}. \quad (3.2)$$

Now, if C is any $n \times n$ matrix, and $\tilde{\mathbf{v}}$ denotes an n -vector of independent *standard* normal random variables, then the quadratic form $\tilde{\mathbf{v}}^\top C \tilde{\mathbf{v}}$ has mean

$$\mathbb{E}\{\tilde{\mathbf{v}}^\top C \tilde{\mathbf{v}}\} = \text{Tr}\{C\}. \quad (3.3)$$

Furthermore, for *symmetric* C , since $\text{Cov}\{\tilde{\mathbf{v}}_i \tilde{\mathbf{v}}_j, \tilde{\mathbf{v}}_k \tilde{\mathbf{v}}_l\} = 0$ unless either $i \neq j$ and $\{i, j\} = \{k, l\}$ or $i = j = k = l$, the variance of $\tilde{\mathbf{v}}^\top C \tilde{\mathbf{v}}$ is given by

$$\text{Var}\{\tilde{\mathbf{v}}^\top C \tilde{\mathbf{v}}\} = \sum_{i=1}^n \sum_{j \neq i}^n C_{ij}(C_{ij} + C_{ji}) + \sum_{i=1}^n C_{ii}^2(\mathbb{E}\{\tilde{\mathbf{v}}_1^4\} - 1) \quad (3.4)$$

$$= 2\text{Tr}(C^2), \quad (3.5)$$

because $\mathbb{E}\{\tilde{\mathbf{v}}_1^4\} = 3$ when $\tilde{\mathbf{v}}_1 \sim \mathcal{N}(0, 1)$; see [16, Chapter 2, Theorem 4 (i) and (ii)]. If C is not symmetric, the expression (3.5) can be used with C replaced by the symmetric matrix $C' := \frac{1}{2}(C + C^\top)$, since $\tilde{\mathbf{v}}^\top C' \tilde{\mathbf{v}} = \tilde{\mathbf{v}}^\top C \tilde{\mathbf{v}}$. Hence, taking expectations on the model with parameters $\theta_0 := (\beta_0, \sigma_0^2, \rho_0)$, we have

$$\mathbb{E}_{\theta_0} s^{(\rho)}(\theta_0; \mathbf{y}) = \text{Tr}\{Z_{\rho_0}\} - \text{Tr}\{Z_{\rho_0}\} = 0,$$

as it must, and

$$\text{Var}_{\theta_0} s^{(\rho)}(\theta_0; \mathbf{y}) = 2\text{Tr}\left\{\left(\frac{1}{2}(Z_{\rho_0} + Z_{\rho_0}^\top)\right)^2\right\} = \text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0} Z_{\rho_0}^\top\}. \quad (3.6)$$

This, together with [17, Theorem 6.2] and [18, Theorem 6.6 and (6.25)], implies the following limit on the precision of estimators of ρ .

Cramér–Rao lower bound. The variance of any unbiased estimator of ρ is no smaller than

$$\gamma_{\rho_0}^2 := \frac{1}{\text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0} Z_{\rho_0}^\top\}}, \quad (3.7)$$

where $Z_\rho := WK_\rho^{-1}$ and K_ρ is as in (2.5).

In practice, the bound implies a precision that cannot be significantly improved upon for any estimator of ρ , without introducing substantial bias; even the best ‘sensible’ estimator of ρ has to have random variability on the scale of γ_{ρ_0} .

Note that, for any $n \times n$ matrix M ,

$$|\text{Tr}\{M^2\}| = \left| \sum_{i=1}^n \sum_{j=1}^n M_{ij} M_{ji} \right| \leq \sum_{i=1}^n \sum_{j=1}^n M_{ij}^2 = \text{Tr}\{MM^\top\}, \quad (3.8)$$

by Cauchy–Schwarz, so that the two traces appearing in $\gamma_{\rho_0}^2$ are not in general the same, and that

$$\text{Tr}\{M^2 + MM^\top\} = \sum_{i=1}^n \sum_{j=1}^n \{M_{ij} M_{ji} + M_{ij}^2\} = \frac{1}{2} \sum_{i=1}^n \sum_{j=1}^n (M_{ij} + M_{ji})^2 \geq 0. \quad (3.9)$$

Consequences of (3.7): the case $\rho_0 = 0$.

As a simple example of what is to be expected, suppose that $\rho_0 = 0$, and that the underlying graph has no isolated vertices, so that $\deg(i) \geq 1$ for all i . Let $W := D^{-1}A$, where D is the diagonal matrix with elements $D_{ii} := \deg(i)$, $1 \leq i \leq n$. Then $Z_{\rho_0} = Z_0 = W$, and we can be specific about the traces in (3.7):

$$\begin{aligned} \text{Tr}\{W^2\} &= \sum_{i=1}^n \sum_{j=1}^n \frac{A_{ij}}{\deg(i) \deg(j)} \\ &\leq \text{Tr}\{WW^\top\} = \sum_{i=1}^n \sum_{j=1}^n \frac{A_{ij}}{\deg(i)^2} = \sum_{i=1}^n \frac{1}{\deg(i)} \leq n. \end{aligned} \quad (3.10)$$

It follows from the equality in (3.10) that $\text{Tr}\{WW^\top\} \geq n/\bar{d}$, where $\bar{d} := n^{-1} \sum_{i=1}^n \deg(i)$ denotes the average degree of a vertex, and hence that $\gamma_0^2 \leq \bar{d}/n$. However, if the vertex degrees are reasonably homogeneous, in the sense all the vertices have degrees at least $c\bar{d}$, for some $0 < c < 1$, then it follows from (3.10) that $\text{Tr}\{WW^\top\} \leq n/c\bar{d}$, and thus, from (3.7), that $\gamma_0^2 \geq c\bar{d}/n$. Hence, in these simple circumstances, the variability of any sensible estimator of ρ is on a scale of at least $(\bar{d}/n)^{1/2}$.

These calculations have particular implications for dense networks. If $W = D^{-1}A$, and if all vertices have degrees exceeding cn , for some fixed $c > 0$, it follows that $\gamma_0 \geq \sqrt{c}$, and it is not possible to obtain an accurate estimate of ρ by any method, no matter how large the size n of the network, if in fact $\rho_0 = 0$. For more general choices of W , $\text{Tr}\{WW^\top\} = \sum_{i,j=1}^n W_{ij}^2$, and if all the elements of W are smaller than C/n , for some $C < \infty$, it follows that $\text{Tr}\{WW^\top\} \leq C^2$, and hence, from (3.7), that $\gamma_0^2 > 1/C^2$. Hence, here too, any estimator of ρ has significant variability on the scale of $1/C$, whatever the size n of the network, if in fact $\rho_0 = 0$.

Consequences of (3.7): general ρ_0 .

If $\rho_0 \neq 0$, analogous results are true. Supposing that $\sum_{j=1}^n W_{ij} = 1$ for each i , it follows that, for all $l \geq 1$, $\sum_{j=1}^n (W^l)_{ij} = 1$ for all i also. Hence, $(W^l)_{ij} = \sum_{k=1}^n (W^{l-1})_{ik} W_{kj}$ is an average of the values W_{1j}, \dots, W_{nj} . Thus, if all the elements of W are smaller than C/n , it follows that $(W^l)_{ij} \leq C/n$ also. Hence it is immediate that, for $l, l' \geq 1$,

$$\text{Tr}\{W^l\} \leq C \quad \text{and} \quad \text{Tr}\{W^l(W^\top)^{l'}\} \leq C. \quad (3.11)$$

Thus, for $\rho \geq 0$, we have

$$\begin{aligned} \text{Tr}\{Z_\rho^2\} &\leq \text{Tr}\{Z_\rho Z_\rho^\top\} = \text{Tr}\{W(I - \rho W)^{-1}(I - \rho W^\top)^{-1}W^\top\} \\ &\leq C \sum_{l \geq 0} (l+1)\rho^l = \frac{C}{(1-\rho)^2}, \end{aligned} \quad (3.12)$$

and so

$$\gamma_{\rho_0} \geq C^{-1/2}(1-\rho).$$

Once again, there is substantial variability in the estimate of ρ unless $C := n \max_{i,j} W_{ij}$ is large. If $\rho < 0$, the bounds in (3.11) can be used to justify similar conclusions, by bounding the difference $\text{Tr}\{Z_\rho Z_\rho^\top\} - \text{Tr}\{WW^\top\}$, provided that $|\rho|$ is not too large.

In particular, in uniformly dense networks, if $W = D^{-1}A$, there is substantial variability in any estimator of ρ if all vertex degrees exceed n/C , and C is not large.

4. Maximum likelihood estimation

The discussion in the previous section shows that there is an upper limit on the precision of any (unbiased) estimator of ρ , and that, in uniformly dense networks, any such estimator must have substantial variability. In this section, we go into more detail about the properties of the MLE.

4.1 General theory

An equation for the MLE

To minimize the quantity $F(\rho; y)$ of (2.14), it is usual to look for values of ρ such that

$$U(\rho; y) := \frac{d}{d\rho} F(\rho; y) = 0. \quad (4.1)$$

Since $F(\rho; y)$ is explicitly given as a function of ρ , its derivative can be calculated, giving

$$\begin{aligned} U(\rho; y) &= \frac{1}{\hat{\sigma}_\rho^2(y)} \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} - \frac{2}{n} \frac{d}{d\rho} \log \det K_\rho \\ &= -\frac{1}{n\hat{\sigma}_\rho^2(y)} (K_\rho y)^\top (I - H_\rho)(Z_\rho^\top + Z_\rho)(I - H_\rho)K_\rho y + \frac{2}{n} \sum_{l \geq 1} \rho^l \text{Tr}\{W^{l+1}\}; \end{aligned} \quad (4.2)$$

the details are given in Appendix A.1. The maximum likelihood estimate $\hat{\rho}_{\text{ML}}$ of ρ is now typically obtained by solving $U(\rho; y) = 0$; Example 3 in Section 4.4 shows that this may not always work.

Note that the data enter through a ratio of quadratic forms in y . Under M_{ρ_0} , as observed in (2.11), $\hat{\sigma}_{\rho_0}^2(y) \sim n^{-1}\sigma^2\chi_{n-m}^2$, having mean $\sigma^2(1 - m/n)$ and variance less than $2n^{-1}\sigma^4$, so that, to a reasonable approximation, $\hat{\sigma}_{\rho_0}^2(y)$ can be replaced by σ^2 in (4.2) if M_{ρ_0} is true, provided that n is large. As a result, $U(\rho; y)$ depends on the data y largely through a single quadratic form in y .

Approximating the distribution of $U(\rho; y)$ under M_{ρ_0}

In order to gauge the typical values of $U(\rho; y)$ to be obtained under the true model M_{ρ_0} , we express $U(\rho; y)$ in terms of the (not directly observable) errors \mathbf{v} , rather than in terms of y . Under M_{ρ_0} , $y = X\beta + K_{\rho_0}^{-1}\mathbf{v}$, and it thus follows from (2.14) and (4.2) that

$$\begin{aligned} U(\rho; y) &= \frac{1}{\hat{\sigma}_\rho^2(y)} \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} - \frac{2}{n} \frac{d}{d\rho} \log \det K_\rho \\ &= -\frac{1}{n\hat{\sigma}_\rho^2(K_{\rho_0}^{-1}\mathbf{v})} \mathbf{v}^\top Q(\rho; \rho_0) \mathbf{v} + \frac{2}{n} \sum_{l \geq 1} \rho^l \text{Tr}\{W^{l+1}\}, \end{aligned} \quad (4.3)$$

where

$$Q(\rho; \rho_0) := (K_\rho K_{\rho_0}^{-1})^\top (I - H_\rho)(Z_\rho^\top + Z_{\rho_0})(I - H_\rho)K_\rho K_{\rho_0}^{-1}. \quad (4.4)$$

Note that β does not appear in (4.3) because $(I - H_\rho)^\top = I - H_\rho$ and $(I - H_\rho)K_\rho X = 0$, implying in particular, using (2.8), that $\hat{\sigma}_\rho^2(\mathbf{y} - X\beta) = \hat{\sigma}_\rho^2(\mathbf{y})$. Hence, and because $\hat{\sigma}_\rho^2(K_{\rho_0}^{-1}\mathbf{v})$ is close to σ^2 for large n , the random part of $U(\rho; \mathbf{y})$ is given, to a close approximation, by $-\mathbf{v}^\top Q(\rho; \rho_0)\mathbf{v}/(n\sigma^2)$.

In the remaining argument, we make frequent use of elementary properties of the trace, that $\text{Tr}(AB) = \text{Tr}(BA)$ and that $\text{Tr}(A^\top) = \text{Tr}(A)$, as well as noting that $z^\top Mz = z^\top M^\top z$ for any compatible z and M , and that H_ρ and $I - H_\rho$ are idempotent. The expectation of $\mathbf{v}^\top Q(\rho_0, \rho_0)\mathbf{v}$, using (3.3), is given by

$$\begin{aligned} \mathbb{E}\{\mathbf{v}^\top Q(\rho_0, \rho_0)\mathbf{v}\} &= \sigma^2 \text{Tr}\{Q(\rho_0, \rho_0)\} = 2\sigma^2 \text{Tr}\{(I - H_{\rho_0})Z_{\rho_0}(I - H_{\rho_0})\} \\ &= 2\sigma^2 \text{Tr}\{(I - H_{\rho_0})Z_{\rho_0}\} = 2\sigma^2 (\text{Tr}\{Z_{\rho_0}\} - \text{Tr}\{H_{\rho_0}Z_{\rho_0}\}), \end{aligned} \quad (4.5)$$

because $K_{\rho_0}K_{\rho_0}^{-1} = I$, and using $\text{Tr}(AB) = \text{Tr}(BA)$ and $(I - H_{\rho_0})^2 = I - H_{\rho_0}$. Then, using $\text{Tr}\{W\} = 0$,

$$\text{Tr}\{Z_{\rho_0}\} = \text{Tr}\{W(I - \rho_0 W)^{-1}\} = \sum_{l \geq 1} \rho_0^l \text{Tr}\{W^{l+1}\}. \quad (4.6)$$

Thus, we have

$$-\mathbb{E}\{\sigma^{-2}\mathbf{v}^\top Q(\rho_0; \rho_0)\mathbf{v}\} + 2 \sum_{l \geq 1} \rho_0^l \text{Tr}\{W^{l+1}\} = 2\text{Tr}\{H_{\rho_0}Z_{\rho_0}\}.$$

Hence, from (4.3) and (4.5), neglecting the small variability in $\hat{\sigma}_\rho^2(\mathbf{y})$, the quantity $nU(\rho_0; \mathbf{y})$ under the model M_{ρ_0} satisfies

$$\mathbb{E}_{\rho_0}\{nU(\rho_0; \mathbf{y})\} \approx \mu_0 := \mu(\rho_0) := 2\text{Tr}\{H_{\rho_0}Z_{\rho_0}\}. \quad (4.7)$$

The principal source of variability in $nU(\rho_0; \mathbf{y})$ arises from $\sigma^{-2}\mathbf{v}^\top Q(\rho_0; \rho_0)\mathbf{v}$, whose variance under M_{ρ_0} , from (3.5), is given by $2\text{Tr}\{Q(\rho_0; \rho_0)^2\}$. This yields

$$\begin{aligned} \text{Var}_{\rho_0}(nU(\rho_0; \mathbf{y})) &\approx 2\text{Tr}\{Q(\rho_0; \rho_0)^2\} \\ &= 2\text{Tr}\{(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})\} \\ &=: \tau_0^2 := \tau_0^2(\rho_0). \end{aligned} \quad (4.8)$$

Combining this with (4.7), we reach the following

Conclusion 1: The distribution of $nU(\rho_0; \mathbf{y})$ can be written in the form

$$nU(\rho_0; \mathbf{y}) = \mu_0 + \tau_0 L(\mathbf{y}), \quad (4.9)$$

where $L(\mathbf{y})$ has mean approximately 0 and variance approximately 1.

Consequences for the distribution of the MLE

Using the information above, we can now make approximate deductions about the MLE $\hat{\rho} := \hat{\rho}_{\text{ML}}$, obtained by setting $U(\rho; y)$ equal to zero. We begin by noting that $U(\hat{\rho}, y) = 0$, whereas, as in (4.9), $nU(\rho_0; y) = \mu_0 + \tau_0 L(y)$, where $L(y)$ approximately has mean zero and variance 1. Writing $U'(\rho; y) := \frac{d}{d\rho} U(\rho; y)$, this implies that

$$-n^{-1}(\mu_0 + \tau_0 L(y)) = \{U(\hat{\rho}; y) - U(\rho_0; y)\} = (\hat{\rho} - \rho_0)U'(\rho_0; y) + o(|\hat{\rho} - \rho_0|),$$

where the final error term is relatively small if $(|\mu_0| + \tau_0)/n|U'(\rho_0; y)|$ is small enough, leading to the approximate expression for the error $\hat{\rho} - \rho_0$ in estimating ρ_0 :

$$\hat{\rho}_{\text{ML}} - \rho_0 \approx -\frac{\mu_0 + \tau_0 L(y)}{nU'(\rho_0; y)}. \quad (4.10)$$

This is essentially a one-step Newton–Raphson approximation to the value of ρ that solves $U(\rho; y) = 0$, starting from the guess ρ_0 . Thus, it is necessary also to examine the magnitude of $nU'(\rho_0; y)$. Since $U(\rho; y)$ is explicitly given, there are exact formulae for its derivative, given in Appendix A.2; however, they are not easy to understand. Instead, it is possible to approximate the *expectation* of $nU'(\rho_0; y)$ on the model M_{ρ_0} , giving

$$n\mathbb{E}_{\rho_0} U'(\rho_0; y) \approx \Delta_0, \quad (4.11)$$

where

$$\Delta_0 = \frac{1}{2}\tau_0^2 + 2\text{Tr}\{-Z_{\rho_0}H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0}) + H_{\rho_0}Z_{\rho_0}^2\}; \quad (4.12)$$

the details are again to be found in Appendix A.2. Thus, from (4.10), we have the following

Conclusion 2: The difference between $\hat{\rho}_{\text{ML}}$ and ρ_0 can be expected to be of magnitude roughly

$$\frac{|\mu_0| \pm \tau_0}{\Delta_0}, \quad (4.13)$$

with the term $|\mu_0|/\Delta_0$ corresponding to systematic bias, and $\psi_0 := \tau_0/\Delta_0$ reflecting the scale of the randomness in the estimator.

Note that the measure ψ_0 should be treated only as a rough guide to the actual variability of $\hat{\rho}_{\text{ML}}$ under M_{ρ_0} , since (4.10) is only a linear approximation to the solution of a non-linear equation, and because there may be added variability arising from the variation of $nU'(\rho_0; y)$ about its mean.

4.2 Estimating ρ when there is no structural element

To get an idea of the relative magnitudes of the terms μ_0 , τ_0 and Δ_0 in (4.13), it is helpful to look at some special cases. For the first, suppose that there is no structural element at all—not even a mean to be fitted—which corresponds to taking $X = 0$; in this case, the formulae (4.7), (4.8) and (4.12) should be interpreted as having $H_\rho = 0$, the projection onto the space $\{0\}$. Then $\mu_0 = 0$, so that any bias is small,

$$\tau_0^2 = 4\text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top\}, \quad \text{where } Z_\rho = W(I - \rho W)^{-1},$$

and

$$\Delta_0 = \tau_0^2/2,$$

giving

$$\psi_0 = \frac{\tau_0}{\Delta_0} = \frac{1}{\sqrt{\text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top\}}}. \quad (4.14)$$

This is the same as γ_{ρ_0} in (3.7). Thus, from (4.13), the typical magnitude of $|\hat{\rho}_{\text{ML}} - \rho_0|$ is of order γ_{ρ_0} . This is the best order of error that could theoretically have been hoped for. Nonetheless, as observed above, it may not lead to accurate estimation of ρ , for instance if $W = D^{-1}A$ and the underlying network is sufficiently dense, even though the *bias* is small.

4.3 Estimating ρ when structure is present

To accommodate the effect of having structure in the model, suppose first that only an overall mean is to be fitted, so that X is the n -vector $\mathbf{1}$, having all elements equal to 1. Suppose also that there are no isolated vertices, and that W is chosen to have $\sum_{j=1}^n W_{ij} = 1$ for all i , so that $W\mathbf{1} = \mathbf{1}$ and the spectral radius $r(W) = 1$. Then it is immediately verified that $H_\rho = H_* := n^{-1}\mathbf{1}\mathbf{1}^\top$ for all ρ , because, for instance, $K_\rho X = (I - \rho W)\mathbf{1} = (1 - \rho)\mathbf{1}$. Hence $WH_\rho = H_*$ and $Z_\rho H_\rho = (1 - \rho)^{-1}H_*$. As a result, by arguments that are made precise in Appendix A.3, (4.8) and (4.12) yield

$$\begin{aligned} \tau_0^2 &= 4\text{Tr}\{(Z_{\rho_0}Z_{\rho_0}^\top + Z_{\rho_0}^2)(I - H_*)\} + 2\{(1 - \rho_0)^{-2} - n^{-1}|Z_{\rho_0}^\top\mathbf{1}|^2\} \\ &\leq 4\text{Tr}\{(Z_{\rho_0}Z_{\rho_0}^\top + Z_{\rho_0}^2)(I - H_*)\}; \\ \Delta_0 &= \frac{1}{2}\tau_0^2 + 2(1 - \rho_0)^{-2}; \end{aligned} \quad (4.15)$$

here, and in what follows, $|x|$ for a vector x denotes the Euclidean norm. Hence, if ρ_0 is fixed and less than 1, ψ_0 is of order $1/(\tau_0 + 1)$, and this, from (4.13), represents the scale of the *random* variation in the estimate of ρ_0 . Now $1/(\tau_0 + 1)$ is small if τ_0^2 is large compared to 1, but can be significantly larger than the lower bound γ_{ρ_0} in (3.7), if τ_0^2 is much smaller than $1/\gamma_{\rho_0}^2 = \text{Tr}\{Z_{\rho_0}Z_{\rho_0}^\top + Z_{\rho_0}^2\}$. As in the case where there is no structure, if $W = D^{-1}A$ and the graph is sufficiently dense, then γ_{ρ_0} need not be small, in which case the MLE has substantial variability.

In contrast to the situation without structure, the principal element μ_0 of the bias in the estimation of ρ need not be zero. If $X = \mathbf{1}$, it follows from (4.7) that

$$\mu_0 = 2\text{Tr}\{Z_{\rho_0}H_0\} = 2(1 - \rho_0)^{-1} > 0.$$

In view of Conclusion 2, the bias becomes important when Δ_0 is not large, which, from (4.15), is when τ_0^2 is not large; as discussed above, this occurs for example if $W = D^{-1}A$ and the graph is uniformly dense. Note also that, from (4.10) and (4.13), the value of the difference $\hat{\rho}_{\text{ML}} - \rho_0$ resulting from μ_0 is roughly $-2/\{\Delta_0(1 - \rho_0)\}$, which typically implies a negative bias, as demonstrated in the results of the simulations referred to in the introduction.

If there is more structure in the regression, but still assuming that the first column of X is $\mathbf{1}$ and that $W\mathbf{1} = \mathbf{1}$, the detailed formulae become more complicated, but the basic message remains the same; the details are given in Appendix A.3. The matrix H_ρ can be represented in the form

$$H_\rho = H_* + \sum_{l=2}^m x_\rho^{(l)} (x_\rho^{(l)})^\top, \quad (4.16)$$

where $H_* := n^{-1}\mathbf{1}\mathbf{1}^\top$ is as before, and $x_\rho^{(2)}, \dots, x_\rho^{(m)}$ are *orthonormal* vectors orthogonal to $\mathbf{1}$ that, with $\mathbf{1}$, span the space generated by the columns of $K_\rho X$. The expression for τ_0^2 derived in (A.13), in addition to the quantity

$$v(\rho_0) := \text{Tr}\{(Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2)(I - H_{\rho_0})\}, \quad (4.17)$$

also involves the traces of the matrices $Z_{\rho_0}^2 H_{\rho_0}$, $Z_{\rho_0}^\top Z_{\rho_0} H_{\rho_0}$, $(Z_{\rho_0} H_{\rho_0})^2$ and $Z_{\rho_0}^\top H_{\rho_0} Z_{\rho_0} H_{\rho_0}$. Similarly, from (4.12), the quantity Δ_0 differs from $\frac{1}{2}\tau_0^2$ by a quantity that is also expressed in terms of these traces. It is shown in Appendix A.3 that if, for instance, it is assumed that each of the vectors $x_\rho^{(l)}$, $2 \leq l \leq m$, has components bounded in modulus by $cn^{-1/2}$, for some fixed $c \geq 1$ —as is the case for $x^{(1)} = n^{-1/2}\mathbf{1}$ —then each of the traces listed above is bounded in modulus by $m^2 c^4 / (1 - |\rho|)^2$. This gives rise to the following

Conclusion 3: If, in the setting of this section, m and c are also fixed for all n , and $|\rho_0|$ is bounded away from 1, then the formulae for τ_0^2 and Δ_0 given in (4.8) and (4.12) differ from $4v(\rho_0)$ and $2v(\rho_0)$ by quantities that are uniformly bounded in n as $n \rightarrow \infty$, where $v(\rho_0)$ is as defined in (4.17). The variability in the MLE of ρ is then on the scale $1/\max\{1, \sqrt{v(\rho_0)}\}$.

Again, this scale of variation may be much larger than γ_0 if $v(\rho_0)$ is much smaller than $\text{Tr}\{Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2\}$.

However, the principal difference is in the quantity μ_0 that is responsible for the bias in estimating ρ , which, from (4.7), (4.10), (4.11) and (A.16), and because $(I - \rho W)^{-1}\mathbf{1} = (1 - \rho)^{-1}\mathbf{1}$, can be expressed in the following

Bias formula: assuming that $W\mathbf{1} = \mathbf{1}$, and that $x_\rho^{(2)}, \dots, x_\rho^{(m)}$ are the remaining vectors in the expression (4.16) for H_ρ , the bias in the MLE is approximately given by $-\mu_0/\Delta_0$, where

$$\mu_0 = 2\text{Tr}\{Z_{\rho_0} H_0\} = 2(1 - \rho_0)^{-1} + \sum_{l=2}^m (x_{\rho_0}^{(l)})^\top Z_{\rho_0} x_{\rho_0}^{(l)}.$$

In the simulations in [5], the remaining columns of X were realized from independent standard normally distributed random variables, and therefore had no particular relationship with A or W , or with each other. In such circumstances, they make no appreciable contribution to the bias. To see this, consider, for simplicity, what happens if $\rho_0 = 0$. Then, for $2 \leq l \leq m$, the *orthonormal* vectors $x_0^{(l)}$ can be thought of roughly as being columns of X , divided by $|x^{(l)}| \approx \sqrt{n}$ so as to have norm 1; ‘roughly’ involves neglecting the small adjustments that are needed, so as to be made orthogonal to $\mathbf{1}$ and to each other. As a result, using (3.3) and (3.5) (with $C := \frac{1}{2}(Z_{\rho_0} + Z_{\rho_0}^\top)$), it follows that typical values of $(x_0^{(l)})^\top Z_0 x_0^{(l)} = (x_0^{(l)})^\top W x_0^{(l)}$ are roughly centred around $n^{-1}\text{Tr} W = 0$, with standard deviation $n^{-1}(\text{Tr}\{\frac{1}{2}(W + W^\top)^2\})^{1/2}$. Hence, they are individually negligible when compared to $\tau_0 \asymp \{\text{Tr}\{WW^\top\}\}^{1/2}$, and thus have no appreciable effect on

the bias. The same is true for their sum, if m is not large. A similar, but more complicated, argument leads to the same conclusion for $\rho_0 \neq 0$. This phenomenon is confirmed in the simulations of [5]. However, if the underlying graph encodes some community structure, with more edges between members of the same community than outside, and if the regression is designed to investigate community related effects, then the resulting contributions to μ_0 can be as important as that from fitting an overall mean, and can have a corresponding effect on the bias. Model $M2$ of Section 5.2 is of this form.

4.4 Ill-conditioned graphs

If the weight matrix associated with the underlying network is ill conditioned, there are consequences for estimation, both of ρ and of β . Taking β first, the distribution of $\hat{\beta}_{\rho_0}$ from (2.8) on the model M_{ρ_0} has covariance matrix $\sigma^2(X^\top S_{\rho_0} X)^{-1}$, as in (2.11), where $S_\rho := K_\rho^\top K_\rho$. It is shown in Appendix A.4 that the variance of the linear combination $a^\top \beta$, for any given m -dimensional unit vector a , is bounded above by

$$\frac{\sigma^2}{\lambda_{\min}(X^\top S_{\rho_0} X)} \leq \frac{\sigma^2}{\lambda_{\min}(X^\top X) \lambda_{\min}(S_{\rho_0})},$$

where $\lambda_{\min}(M)$ and $\lambda_{\max}(M)$ denote the smallest and largest eigenvalues of a symmetric matrix M . Thus, for fixed σ^2 , provided that $\lambda_{\min}(S_{\rho_0})$ is not close to zero, linear combinations of β can be accurately estimated if $\lambda_{\min}(X^\top X)$ is large, which is the usual condition for ordinary least squares. And, even if $\lambda_{\min}(S_{\rho_0})$ is small, it may still be the case that $\lambda_{\min}(X^\top S_{\rho_0} X)$ is comparable to $\lambda_{\min}(X^\top X)$, if the space spanned by the columns of X does not contain vectors that are substantially shrunk by $S_{\rho_0}^{1/2}$.

If ρ_0 is not accurately estimated, it will typically be the case that $\hat{\beta}_\rho$ is used to estimate β , for some value $\rho \neq \rho_0$. In such cases, it is shown in Appendix A.4 that linear combinations of β can be estimated using $\hat{\beta}_\rho$ to an accuracy comparable to that obtained using $\hat{\beta}_{\rho_0}$, provided that $\lambda_{\max}(S(\rho, \rho_0))$ is not large, where $S(\rho, \rho_0) := S_\rho^{1/2} S_{\rho_0}^{-1} S_\rho^{1/2}$. However, if $\lambda_{\max}(S(\rho, \rho'))$ can be large for values of ρ and ρ' that are plausible, estimation of β may be rather less accurate. In particular, it is shown in Appendix A.4 that a large value of $\lambda_{\max}(W^\top W)$ indicates that problems with the estimation of β may also arise; see Example 1 in Section 4.4.

For bias and variability in the estimation of ρ , it has already been shown that problems can easily arise in dense networks, without W having any unusual structure. However, these problems can be accentuated by particular networks and design matrices. Again assuming that $W\mathbf{1} = \mathbf{1}$, this is because the trace $v(\rho_0) = \text{Tr}\{(Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2)(I - H_{\rho_0})\}$ in (4.17) can be very much smaller than $\text{Tr}\{Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2\} = 1/\gamma_{\rho_0}^2$. If $\tau_0 \geq 1$, so that $\psi_0 := \tau_0/\Delta_0 \asymp 1/(\tau_0 + 1)$, this leads to a value of ψ_0 that is much larger than the theoretical lower bound γ_{ρ_0} . Taking $\rho_0 = 0$ for simplicity, the comparison is between $\text{Tr}\{(WW^\top + W^2)(I - H_0)\}$ and $\text{Tr}\{WW^\top\}$.

Example 1. As a first example, take that of the $(n-1)$ -star, for which $A_{ij} = A_{ji} = 1$, $2 \leq j \leq n$, and $A_{ij} = 0$ otherwise, and take $W := D^{-1}A$. Then the matrix W^2 has $(W^2)_{11} = 1$, and all other elements less than or equal to $1/(n-1)$. If X consists only of the single column $\mathbf{1}$, then $H_0 := H_* := n^{-1}\mathbf{1}\mathbf{1}^\top$, $WH_0 = H_*$, and the matrix $W^2H_0 = H_*$ has all its elements equal to $1/n$. Thus the traces of W^2 and W^2H_0 are bounded for all n ; direct calculation gives $\text{Tr}\{W^2\} = 2$ and $\text{Tr}\{W^2H_0\} = 1$. The matrix WW^\top consists of an $(n-1) \times (n-1)$ block of 1's, for $2 \leq i, j \leq n$, together with $1/(n-1)$ at position $(1, 1)$, and the remaining elements zero; the matrix $WW^\top H$ has a corresponding $(n-1) \times (n-1)$ block, now with values $1 - 1/n$, all elements in the first row have the value $1/n(n-1)$, and the remaining elements in the first column have value $1 - 1/n$. Thus the diagonal elements of $WW^\top(I - H_0)$ all have values equal

to $1/n$, and so its trace is exactly 1, whereas $\text{Tr}\{WW^\top\} = n - 1 + 1/(n - 1) \geq n - 1$. Hence the lower bound γ_0 at $\rho = 0$ is less than $1/(n - 1)$, but, with the simple structure corresponding to estimating the overall mean in the model, using the formula (4.15) from Section 4.3,

$$\tau_0^2 = 4\text{Tr}\{(WW^\top + W^2)(I - H_*)\} = 4(1 + 2 - 1) = 8$$

is comparable to 1, implying that the MLE does not give accurate estimation of ρ .

For estimating β , note that, from (2.8) and (2.11), the variance of $\hat{\beta}_\rho$ on the model M_ρ is given by $\sigma^2(\mathbf{1}^\top K_\rho^\top K_\rho \mathbf{1})^{-1} = \sigma^2/\{n(1 - \rho)^2\}$, because $K_\rho \mathbf{1} = (1 - \rho)\mathbf{1}$. Hence, if ρ is known, β can be estimated accurately by $\hat{\beta}_\rho$ if n is large and ρ is not too close to 1. However, in this example, if $e^{(1)}$ denotes the vector $(1, 0, \dots, 0)^\top$, then $\lambda_{\max}(W^\top W) \geq (e^{(1)})^\top W^\top W e^{(1)} = n - 1$ is large, suggesting that using the wrong value of ρ in estimating β may cause problems. Applying (A.22) from Appendix A.4 below, and taking $\rho_0 = 0$ to simplify the calculations, since then $S_{\rho_0} = I$, we obtain the explicit expression

$$\begin{aligned} \text{Var}_0 \hat{\beta}_\rho &= \frac{\sigma^2}{n^2(1 - \rho)^2} \mathbf{1}^\top K_\rho K_\rho^\top \mathbf{1} = \frac{\sigma^2}{n^2(1 - \rho)^2} \{n(1 - 2\rho) + \rho^2 \mathbf{1}^\top WW^\top \mathbf{1}\} \\ &= \frac{\sigma^2}{n^2(1 - \rho)^2} \left\{ n(1 - 2\rho) + \rho^2 \left((n - 1)^2 + \frac{1}{n - 1} \right) \right\}. \end{aligned}$$

Hence, the variance of $\hat{\beta}_\rho$ grows like $\sigma^2\{\rho/(1 - \rho)\}^2$ away from $\rho_0 = 0$ and is thus substantial when the wrong value of ρ is used. As a result, since ρ cannot be reliably estimated, nor can β .

Conclusion 4: This example shows that having a large number of vertices of low degree—in this case, $n - 1$ vertices of degree 1—is *not* of itself enough to imply accurate estimation of ρ or β , even though it ensures that $\text{Tr}\{WW^\top\}$ is large, and hence that γ_0 is small.

Adding further columns to X results in a projection matrix $H = H_* + \tilde{H}$, with $H_*\tilde{H} = 0$, $\tilde{H}^\top = \tilde{H}$ and $\tilde{H}^2 = \tilde{H}$, replacing the trace $\text{Tr}\{WW^\top(I - H_*)\}$ by $\text{Tr}\{WW^\top(I - H)\}$; the latter is smaller, because

$$\begin{aligned} &\text{Tr}\{WW^\top(I - H_*)\} - \text{Tr}\{WW^\top(I - H)\} \\ &= \text{Tr}\{WW^\top(I - H_*)\} - \text{Tr}\{(I - H_* - \tilde{H})WW^\top(I - H_* - \tilde{H})\} \\ &= \text{Tr}\{WW^\top(I - H_*)\} - \text{Tr}\{(I - H_*)WW^\top(I - H_*) - 2\tilde{H}WW^\top(I - H_*) + \tilde{H}WW^\top\tilde{H}\} \\ &= \text{Tr}\{2(I - H_*)\tilde{H}WW^\top - \tilde{H}WW^\top\} = \text{Tr}\{\tilde{H}WW^\top\} \\ &= \text{Tr}\{(\tilde{H}W)(\tilde{H}W)^\top\} \geq 0. \end{aligned} \tag{4.18}$$

The value of $\text{Tr}\{W^2H\}$ is not materially changed, because, apart from $W_{11} = 1$, the elements of W^2 are uniformly of order $O(1/n)$. Hence, even with the more realistic model structure, the value of $v(\rho_0)$ is still comparable to 1 when $\rho_0 = 0$, and, in view of Conclusion 3, the variability of the MLE of ρ is substantial.

Example 2. The same sort of discussion can be carried out, if the network above is generalized to a small number of stars, each of which has high degree, the remaining vertices being leaves, and the centres of the stars being connected among themselves. The elements of W^2 , apart from the diagonal elements corresponding to the centres of the stars, are then uniformly small, as in (3.10), typically resulting in only moderate values of $\text{Tr}\{W^2\}$ and $\text{Tr}\{W^2H\}$. The block structure of WW^\top in the previous example is

replaced by a collection of blocks of 1's, one for the leaves of each star. As a result, by arguments similar to those in the example above, the trace $\text{Tr}\{WW^\top(I-H)\}$ becomes small if the design matrix X contains columns identifying the blocks; the elements of the column identifying a block should take the value 1 at positions within the corresponding block, and zero within the other blocks, the remaining elements being chosen to ensure that the sum of the columns is $\mathbf{1}$, so that the overall mean belongs to the design. In consequence, the value of $v(0)$ is comparable to 1, and, in view of Conclusion 3, the variability of the MLE of ρ is substantial in this case as well.

Example 3. A further, more extreme example, is the one considered by [7], in which $A = \mathbf{1}\mathbf{1}^\top - I$ is the adjacency matrix of the *complete* graph on $n \geq 2$ vertices. Then, for $W = D^{-1}A$, we have

$$W = W^\top = \frac{1}{n-1}\{\mathbf{1}\mathbf{1}^\top - I\} = H_* - \frac{1}{n-1}\{I - H_*\}, \quad (4.19)$$

where, as before, $H_* = n^{-1}\mathbf{1}\mathbf{1}^\top$, implying, since H_* and $I - H_*$ are idempotent and $H_*(I - H_*) = 0$, that

$$W^r = H_* + \left(\frac{-1}{n-1}\right)^r \{I - H_*\} \quad \text{and} \quad Z_\rho = \frac{H_*}{1-\rho} - \frac{I - H_*}{n-1+\rho} = Z_\rho^\top. \quad (4.20)$$

Hence, in particular,

$$\text{Tr}\{Z_\rho Z_\rho^\top + Z_\rho^2\} = 2\left\{\frac{1}{(1-\rho)^2} + \frac{n-1}{(n-1+\rho)^2}\right\},$$

and so, in view of (3.7), estimation of ρ cannot be accurate if ρ is not close to 1. However, if the model consists only of an intercept term, so that $X = \mathbf{1}$, and recalling that $\mathbf{1}^\top W = \mathbf{1}^\top$, the estimate of β from (2.8) is given by

$$\hat{\beta}_\rho(y) = (\mathbf{1}^\top(I - \rho W)^2\mathbf{1})^{-1}\mathbf{1}^\top(I - \rho W)^2y = n^{-1}\sum_{j=1}^n y_j,$$

which is the same for all values of ρ , and its variance, obtained using the general formula in (A.22) below, reduces to $1/\{n(1 - \rho_0)^2\}$ on M_{ρ_0} , and becomes small as n increases. Thus, although ρ may not be accurately estimable, there is no problem in estimating β .

Returning to the estimation of ρ , still with $X = \mathbf{1}$, note also that, using (4.20),

$$\text{Tr}\{(I - H_0)WW^\top\} = \text{Tr}\{(I - H_*)W^2\} = \frac{1}{(n-1)^2}\text{Tr}\{I - H_*\} = \frac{1}{n-1}.$$

Thus, $\text{Tr}\{(I - H_0)WW^\top\}$ is an order of magnitude smaller than $\text{Tr}\{WW^\top\}$. This in turn, from (4.15), implies that the variability in the MLE of ρ is on the scale of $\tau_0 \ll 1$, in apparent contradiction to the lower bound in (3.7); this can only be the case if the *bias* in the MLE is extremely large. We now investigate what happens in more detail.

First, allowing a more general design matrix X with m columns, observe that

$$WX = (n-1)^{-1}\{\mathbf{1}\mathbf{1}^\top X - X\},$$

so that the i th column of WX is a linear combination of $\mathbf{1}$ and the i th column of X . If X contains the column $\mathbf{1}$, this implies that the columns of $K_\rho X = X - \rho WX$ span the same space as the columns of X , and hence that the orthogonal projection H_ρ is the same for all ρ : $H_\rho = H_0 = H(X)$, and $H_0 \mathbf{1} = \mathbf{1}$. As a result, it follows that

$$H_0 H_* = H_* H_0 = H_* \quad \text{and} \quad (I - H_0) H_* = H_* (I - H_0) = 0, \quad (4.21)$$

so that, from (4.8), (4.20) and (4.21),

$$\tau_0^2 = 8 \text{Tr}\{Z_{\rho_0}(I - H_0)Z_{\rho_0}(I - H_0)\} = \frac{8}{(n-1+\rho_0)^2} \text{Tr}\{I - H_0\} = \frac{8(n-m)}{(n-1+\rho_0)^2} \approx \frac{8}{n}.$$

Similar calculations show that, using (4.20) and (4.21) in (4.12), there is great simplification, yielding

$$\begin{aligned} \Delta_0 &= \frac{1}{2} \tau_0^2 + 2 \text{Tr}\{-Z_{\rho_0} H_0 (Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_0) + H_0 Z_{\rho_0}^2\} \\ &= \frac{1}{2} \tau_0^2 + 2 \left\{ \frac{1}{(1-\rho_0)^2} + \frac{m-1}{(n-1+\rho_0)^2} \right\} = \frac{2}{(1-\rho_0)^2} + O(n^{-1}), \end{aligned}$$

since $\text{Tr}\{H_0 - H_*\} = m-1$; and then

$$\mu_0 = 2 \text{Tr}\{H_0 Z_{\rho_0}\} = \frac{2}{1-\rho_0} - \frac{m-1}{n-1+\rho_0}.$$

Hence it follows that μ_0/Δ_0 is not small, and that the bias in the MLE is thus substantial.

As it happens, using (4.19) and (4.21),

$$(I - H_\rho) K_\rho y = (I - H_0)(I - \rho W)y = \left\{1 + \frac{\rho}{n-1}\right\} (I - H_0)y,$$

and, from (2.8), and because $(I - H_\rho) = (I - H_\rho)^2 = (I - H_0)^2$, this implies that

$$\hat{\sigma}_\rho^2(y) = n^{-1} (K_\rho y)^\top (I - H_\rho) K_\rho y = \left\{1 + \frac{\rho}{n-1}\right\}^2 n^{-1} y^\top (I - H_0)y.$$

Thus, $\hat{\sigma}_\rho^2(y)$ varies very little as ρ ranges over the interval $(-1, 1)$. What is worse, $\log(\hat{\sigma}_\rho^2(y))$ is the sum of a function of ρ alone and a function of the data alone, as is therefore the function $F(\rho; y)$ in (2.14) also. Thus, in this setting, the maximum likelihood estimator, choosing $\rho \in (-1, 1)$ to minimize $F(\rho; y)$, represents a choice on which the data have no influence, so that its variability is actually zero, and the same fixed choice of ρ would always be made.

Conclusion 5: For the complete graph, the likelihood gives no information at all about the value of ρ .

For this choice of W , K_ρ has eigenvalues $1 - \rho$, together with $1 + \rho/(n-1)$ repeated $n-1$ times, and so the function

$$\begin{aligned} F(\rho; y) &= \log\{n^{-1} y^\top (I - H_0)y\} + 2 \log\{1 + \rho/(n-1)\} \\ &\quad - \frac{2}{n} \{\log(1 - \rho) + (n-1) \log\{1 + \rho/(n-1)\}\} \\ &= \log\{n^{-1} y^\top (I - H_0)y\} + \frac{2}{n} \{\log\{1 + \rho/(n-1)\} - \log(1 - \rho)\} \end{aligned}$$

defined in (2.14) has strictly positive derivative throughout the (open) interval $(-\infty, 1)$. Hence setting its derivative equal to zero to find a minimum would also fail here.

5. Estimation using a quadratic form

The conclusions of Section 4 indicate that, in dense graphs, the MLE of the correlation parameter ρ in the network disturbance model may quite generally exhibit significant bias and variability. The ill-conditioned networks of Section 4.4 indicate that these effects can be dramatic. Because of the Cramér–Rao bound given in (3.7), it is not possible to avoid the variability, whatever estimator of ρ is used. However, estimators with smaller bias can be envisaged; see [8–10] for progress in the context of correcting the bias in the MLE for network effects models, when the variability is small. In this section, we move away from the MLE, and instead propose a simple and intuitive estimator of ρ , which suffers much less from bias, even in circumstances in which all estimators have substantial variability.

5.1 A simple estimator of ρ

An important element in $n \frac{dF(\rho; y)}{d\rho}$, the quantity that is set equal to zero for the MLE, is the quadratic form

$$(K_\rho y)^\top (I - H_\rho)(Z_\rho + Z_\rho^\top)(I - H_\rho)K_\rho y$$

in y , as can be seen in (4.2). Here, we replace this quadratic form with a simpler one. Since, under the model M_{ρ_0} from (2.7), the innovations $\mathbf{v}_i = (K_{\rho_0} \boldsymbol{\epsilon})_i$, $1 \leq i \leq n$, are independent, even at neighbouring vertices, the expectation of $(K_{\rho_0} \boldsymbol{\epsilon})_i (K_{\rho_0} \boldsymbol{\epsilon})_j$ is zero for all $i \neq j$. On the other hand, for $\rho \neq \rho_0$, the expectation of $(K_\rho \boldsymbol{\epsilon})_i (K_\rho \boldsymbol{\epsilon})_j$ under M_{ρ_0} is typically not zero, if i and j are neighbours. For example, if $\rho_0 = 0$, then, for $i \neq j$,

$$\sigma^{-2} \mathbb{E}_0\{(K_\rho \boldsymbol{\epsilon})_i (K_\rho \boldsymbol{\epsilon})_j\} = \{(I - \rho W)(I - \rho W)^\top\}_{ij} = -\rho(W_{ij} + W_{ji}) + \rho^2(WW^\top)_{ij}, \quad (5.1)$$

which is not zero whenever $\rho \neq 0$ is small enough, if one of W_{ij} and W_{ji} is non-zero. This suggests building sums

$$T_\rho^C(y) := (K_\rho \hat{\boldsymbol{\epsilon}}_\rho(y))^\top C K_\rho \hat{\boldsymbol{\epsilon}}_\rho(y) \quad (5.2)$$

of products of the form $(K_\rho \hat{\boldsymbol{\epsilon}}_\rho(y))_i (K_\rho \hat{\boldsymbol{\epsilon}}_\rho(y))_j$, where $\hat{\boldsymbol{\epsilon}}_\rho(y) := K_\rho^{-1}(I - H_\rho)K_\rho y$ is the vector of residuals after fitting M_ρ , as derived in (2.10); thus

$$T_\rho^C(y) = y^\top K_\rho^\top (I - H_\rho) C (I - H_\rho) K_\rho y. \quad (5.3)$$

The matrices C , having non-negative coefficients that do not depend on the data or on ρ , and with all diagonal elements equal to zero, are to be suitably chosen. Any such sum $T_\rho^C(y)$, for fixed C and y , should take values close to zero when ρ is close to ρ_0 , and this can potentially be exploited to gain information about the value of ρ . As indicated in (5.1), the elements of C should be chosen so as to give larger weight to pairs (i, j) where $W_{ij} + W_{ji}$ is large, and a natural choice is thus to take $C = \frac{1}{2}(W + W^\top)$, or, equivalently, to take $C = W$. Another intuitive choice, reflecting the structure of the underlying graph, rather than the model weights, is to take $C = A$, roughly analogous to the lag-one autocovariance in time series analysis.

Recalling that, on M_{ρ_0} , $\mathbf{y} = X\beta + K_{\rho_0}^{-1}\mathbf{v}$, it follows that, on M_{ρ_0} ,

$$\hat{\varepsilon}_\rho(\mathbf{y}) = K_\rho^{-1}(I - H_\rho)K_\rho K_{\rho_0}^{-1}\mathbf{v},$$

since $(I - H_\rho)K_\rho X = 0$, and hence that, on M_{ρ_0} ,

$$T_\rho^C(\mathbf{y}) = \mathbf{v}^\top Q_C(\rho, \rho_0)\mathbf{v}, \quad (5.4)$$

where, much as in (4.4),

$$Q_C(\rho, \rho_0) := (K_\rho K_{\rho_0}^{-1})^\top (I - H_\rho)C(I - H_\rho)K_\rho K_{\rho_0}^{-1}. \quad (5.5)$$

From (5.4) and (5.5), we have

$$\mathbb{E}_{\rho_0} T_{\rho_0}^C(\mathbf{y}) = \sigma^2 \text{Tr}\{(I - H_{\rho_0})C\} = -\sigma^2 \text{Tr}\{H_{\rho_0}C\} =: \sigma^2 \hat{\mu}_{\rho_0}^C, \quad (5.6)$$

since C has zero diagonal, and so $\text{Tr}\{C\} = 0$. Equation (5.6) suggests the following ‘method of moments’ procedure for estimating ρ .

Elementary estimate of ρ : find $\rho = \hat{\rho}_C$ so that

$$T_\rho^C(y) + \hat{\sigma}_\rho^2(y) \text{Tr}\{H_\rho C\} = 0. \quad (5.7)$$

Equation (5.7) is used to correct the bias that would result, in view of (5.6), if, in estimating ρ , $T_\rho^C(y)$ were itself set equal to zero.

Procedure:

Input: a network with weight matrix W , a linear model with design matrix X and observed data y , and a non-negative coefficient matrix C with zero diagonal.

Function: for any value of ρ :

- (1) Calculate the trace $\text{Tr}\{H_\rho C\}$, using (2.9) and (2.2);
- (2) Calculate the parameter estimates $\hat{\beta}_\rho(y)$ and $\hat{\sigma}_\rho^2(y)$ using (2.8);
- (3) Calculate $\hat{\varepsilon}_\rho(y)$, using $\hat{\beta}_\rho(y)$ and (2.10);
- (4) Calculate the value of $T_\rho^C(y)$ from (5.2).

Solution: Use a root finding algorithm to find the zero of $T_\rho^C(y) + \hat{\sigma}_\rho^2(y) \text{Tr}\{H_\rho C\}$, which is the estimate $\hat{\rho}_C$ of the correlation ρ using T^C .

In this article, we do not attempt to give detailed conditions under which the procedure using (5.7) gives good estimates $\hat{\rho}_C$ of ρ . It is clear from the example of the complete graph in Section 4.4 that this cannot always be so; indeed, for the complete graph, (5.7) has no solution. Instead, we carried out some simulations, illustrating that $\hat{\rho}_C$ can be effective, within the limitations of what is possible. These are

discussed in Section 5.2 below. In the remainder of this section, we give a rough guide to its precision from a theoretical standpoint, as well as proposing a permutation procedure for assessing its precision based on the data to hand.

The scale of variation of $\hat{\rho}_C$

To determine the scale of variation about ρ_0 of the resulting estimator $\hat{\rho}_C$, write

$$U_\rho^C(y) := T_\rho^C(y) + \hat{\sigma}_\rho^2(y) \text{Tr}\{H_\rho C\}. \quad (5.8)$$

As in the discussion of the MLE around (4.10) and (4.11) in Section 4, a plausible measure is given by $\hat{\psi}_{\rho_0}^C := \hat{\tau}_{\rho_0}^C / |\hat{\Delta}_{\rho_0}^C|$, to be estimated by $\hat{\tau}_{\hat{\rho}_C}^C / |\hat{\Delta}_{\hat{\rho}_C}^C|$, where

$$(\hat{\tau}_{\rho_0}^C)^2 := \sigma^{-4} \text{Var}_{\rho_0}(U_{\rho_0}^C(\mathbf{y})) \quad \text{and} \quad \hat{\Delta}_{\rho_0}^C := -\sigma^{-2} \mathbb{E}_{\rho_0} \left\{ \frac{dU_\rho^C(\mathbf{y})}{d\rho} \Big|_{\rho=\rho_0} \right\}.$$

The quantity $(\hat{\tau}_{\rho_0}^C)^2$ is dominated by the trace of a matrix,

$$\text{Tr}\{(I - H_{\rho_0})(C + C^\top)(I - H_{\rho_0})(C + C^\top)\}, \quad (5.9)$$

which is reminiscent of the definition of τ_0^2 in (4.8), and the quantity $\hat{\Delta}_{\rho_0}^C$ has leading term

$$\text{Tr}\{(I - H_{\rho_0})[Z_{\rho_0}^\top(I - H_{\rho_0}) - Z_{\rho_0}H_{\rho_0}](C + C^\top)\}. \quad (5.10)$$

The derivations, and more detailed expressions, are given in Section A.5. As before, the scale of variation

$$\hat{\psi}_{\rho_0}^C := \hat{\tau}_{\rho_0}^C / |\hat{\Delta}_{\rho_0}^C|, \quad (5.11)$$

determined from (5.9) and (5.10), need not correspond precisely to the standard deviation of $\hat{\rho}_C$ on M_{ρ_0} , because $U_\rho^C(y)$ is not linear as a function of ρ , and because the derivative of $U_\rho^C(y)$ may fluctuate appreciably about its expectation.

To get an idea of the magnitude of $\hat{\psi}_{\rho_0}^C$, let $C = W$, and suppose that $W\mathbf{1} = \mathbf{1}$ and that $X = \mathbf{1}$, so that only the overall mean is being fitted. Then, as in Section 4.3, $K_\rho X = (1 - \rho)\mathbf{1}$, implying that, for all ρ , $H_\rho = H_* := n^{-1}\mathbf{1}\mathbf{1}^\top = WH_\rho$ and that $Z_\rho H_\rho = (1 - \rho)^{-1}H_*$. As a result, we easily deduce that

$$(\hat{\tau}_{\rho_0}^W)^2 = 2(\text{Tr}\{W^2 + WW^\top - H_*WW^\top\} - 1) + O(n^{-1}),$$

and that

$$\hat{\Delta}_{\rho_0}^C = -\text{Tr}\{WZ_{\rho_0} + WZ_{\rho_0}^\top - H_*WZ_{\rho_0}^\top\} + (1 - \rho)^{-1} + n^{-1}\text{Tr}\{(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_*)\},$$

where we have used the full expressions in (A.25) and (A.30) instead of (5.9) and (5.10) to obtain the smaller order terms. In particular, if $\rho_0 = 0$, so that $Z_{\rho_0} = W$, $\frac{1}{2}(\hat{\tau}_0^W)^2 = -\hat{\Delta}_0^C + O(n^{-1})$, and the magnitude of $\hat{\psi}_0^C$ is roughly $2/\hat{\tau}_0^W$; this remains the case for ρ_0 not too far from zero. As in the discussion in Section 3, this implies that if, for instance, $W_{ij} \leq c/n$ for all i, j , then the scale of variation is at least

$1/\sqrt{c}$, and is appreciable if c is not large. For $W := D^{-1}A$, this corresponds to having a dense underlying graph, as before.

The justification for using (5.7) to estimate ρ relies only on the first two joint moments of the vector \mathbf{v} of errors, and does not assume that the components are independent or normally distributed, though they should be uncorrelated. However, the formula (A.25), used in the theoretical discussion of the scale of random variation $\hat{\psi}_{\rho_0}^C$ of the estimator $\hat{\rho}_C$, is already based on stronger assumptions. For instance, if the errors \mathbf{v} are independent and have zero mean, but if $\mathbb{E}\{\sigma^{-4}\mathbf{v}_1^4\} \neq 3$, the expression (3.4) for the variance of a quadratic form needs to be used in place of (3.5).

A permutation procedure for assessing the accuracy of $\hat{\rho}_C$

Instead of assessing the precision of $\hat{\rho}_C$ using $\hat{\psi}_{\rho_0}^C$, some form of simulation procedure can be used. One possibility is the following permutation algorithm. Start by computing the estimated \mathbf{v} -residuals

$$\hat{\mathbf{v}}(y) := K_{\hat{\rho}_C} \hat{\varepsilon}_{\hat{\rho}_C}(y) = (I - H_{\hat{\rho}_C}) K_{\hat{\rho}_C} y. \quad (5.12)$$

Since the components of \mathbf{v} are independent and identically distributed, this suggests that synthetic data could be derived by replacing the \mathbf{v} -residuals $(\hat{v}_i(y), 1 \leq i \leq n)$ by permuted residuals

$$\hat{v}_i^\pi(y) := \hat{v}_{\pi(i)}(y), \quad 1 \leq i \leq n, \quad (5.13)$$

where π is a randomly chosen permutation of $\{1, 2, \dots, n\}$. These can be used to determine residuals

$$\hat{\varepsilon}^\pi(y) := K_{\hat{\rho}_C}^{-1} \hat{\mathbf{v}}^\pi(y) \quad (5.14)$$

for the original model, and hence to generate synthetic data $y^\pi := X\hat{\beta} + \hat{\varepsilon}^\pi(y)$. These can in turn be analysed in the same way as the original data, to provide an estimate $\hat{\rho}_C^\pi$ of $\hat{\rho}_C(y)$. Repeating this procedure for a large number M of independently chosen random permutations π_j , $1 \leq j \leq M$, yields a selection of estimates $\hat{\rho}_C^{\pi_j}$ of $\hat{\rho}_C(y)$, whose empirical variability should be much like that of the variability of $\hat{\rho}_C$ about the model value ρ_0 , when $\rho_0 = \hat{\rho}_C(y)$.

5.2 Simulations

In this section, we illustrate the effectiveness of $\hat{\rho}_C$ using simulated data, showing that it substantially reduces the bias exhibited by the MLE, while its precision remains constrained by the lower bound (3.7). In Table 1, data were simulated at the vertices of a Bernoulli graph $G(n, p)$, with $n = 100$ and with three choices of p . The data were simulated from the model M_ρ with six different values of ρ . The parameters $\beta = (1, 0.5, 0.4, 0.3)$ were held fixed, and the design matrix X was chosen to have $\mathbf{1}$ as its first column, with its remaining columns consisting of independent realizations of standard normal random variables. For each combination of p and ρ , the elements of X and the network G were sampled anew; then 100 replicates were generated, starting from realizations of independent standard normal random variables \mathbf{v} , and using (1.2) to determine the values of ε , taking $W_{ij} := D_i^{-1}A_{ij}$ if $D_i \geq 1$, and $W_{ij} = 0$ if i is an isolated vertex. Each cell in Table 1(a) contains the mean of the corresponding 100 values of $\hat{\rho}_W$, calculated using (5.7) with $C = W$, together with the standard error of a single estimate. It is clear that the standard errors are substantial. The means are close to the truth, within the limits of variability; for 100 replicates, the standard error of the mean is one tenth of that of a single estimate. The estimates of β were uniformly good, and are not presented. In a Bernoulli graph $G(n, p)$, the degrees of the vertices

TABLE 1 Comparison between $\hat{\rho}_C$ and $\hat{\rho}_{ML}$ for a network disturbance model based on $G(100, p)$ with $W = D^{-1}A$.

		p		
		0.0975	0.19	0.36
ρ	−0.2	−0.21 (0.20)	−0.24 (0.30)	−0.24 (0.44)
	−0.1	−0.14 (0.23)	−0.10 (0.35)	−0.11 (0.46)
	0.0	0.01 (0.23)	−0.03 (0.37)	0.02 (0.47)
	0.1	0.08 (0.24)	0.09 (0.31)	0.08 (0.46)
	0.2	0.21 (0.23)	0.16 (0.34)	0.11 (0.50)
	0.3	0.28 (0.21)	0.23 (0.32)	0.27 (0.47)
CRLB	0.0	(0.20)	(0.29)	(0.39)

(a) Means of 100 simulated estimates of ρ (with the standard error of a single estimate) calculated using (5.7) with $C = W$. The Cramér–Rao lower bound (CRLB) line gives the values of the lower bound $\gamma_{\rho 0}$ derived using (3.7).

		p		
		0.0975	0.19	0.36
ρ	−0.2	−0.26 (0.21)	−0.28 (0.28)	−0.33 (0.39)
	−0.1	−0.13 (0.20)	−0.22 (0.34)	−0.28 (0.37)
	0.0	−0.03 (0.21)	−0.10 (0.30)	−0.19 (0.42)
	0.1	0.02 (0.21)	−0.05 (0.30)	−0.11 (0.41)
	0.2	0.13 (0.20)	0.14 (0.31)	−0.05 (0.42)
	0.3	0.19 (0.20)	0.16 (0.30)	0.02 (0.33)
CRLB	0.0	(0.20)	(0.29)	(0.39)

(b) Means of 100 simulated MLE estimates of ρ . The standard errors of a single observation are consistent with the values given in the CRLB line.

are reasonably homogeneous, being concentrated around their mean $(n - 1)p$ when $p \gg 1/n$. Thus, from (3.10), calculating as if all the degrees were exactly $(n - 1)p$, the lower bound (3.7) on the variance of an estimator when $\rho = 0$ gives γ_ρ approximately equal to $\sqrt{p/2}$. The actual values of γ_ρ in the simulations depend on the value of ρ and on the particular realization of the Bernoulli random graphs that were sampled. The value of ρ made little difference, but the randomness in the degrees was still enough, for $n = 100$, to reduce the value of γ_ρ by about 10%; typical values are quoted in the final line of Table 1. They indicate that the observed variability of the quadratic form estimator is not far from the best that could be achieved in this setting.

For comparison, we give the results in Table 1(b) of using the MLE in the same setting. The standard deviations were uniformly close to the values given in the CRLB line in the table above and were thus a little smaller than those in the experiments using (5.7) with $C = W = D^{-1}A$. However, the MLE exhibits a systematic bias, which becomes progressively worse as the density of the graph increases. Even with edge density $p = 0.0975$, the smallest value investigated, the means of the estimates of ρ were consistently too small, by amounts comparable to 0.05, whereas the whole range of values of ρ being investigated was only 0.5.

In the second experiment, the graph $G(n, p)$ was replaced by an Erdős–Rényi mixture model, consisting of two classes with 50 vertices in each; the within class probabilities were taken to be $2p(1 - p)$ and the

TABLE 2 Comparison between $\hat{\rho}_C$ and $\hat{\rho}_{ML}$ for two network disturbance models based on an Erdős–Rényi mixture graph, each with $W := D^{-1}A$ and with $\rho = 0.1$.

p		0.05	0.1	0.2
0.8	$M1$	0.102 (0.20)	0.104 (0.28)	0.111 (0.43)
	$M2$	0.094 (0.21)	0.082 (0.29)	0.098 (0.42)
	CRLB	(0.18)	(0.26)	(0.35)

(a) Means of 1,000 simulated estimates of ρ (with the standard error of a single estimate) using (5.7) with $C = W$. The CRLB line gives the values of the lower bound γ_{ρ_0} derived using (3.7).

p		0.05	0.1	0.2
0.8	$M1$	0.06 (0.18)	0.01 (0.26)	−0.09 (0.37)
	$M2$	0.05 (0.19)	−0.02 (0.27)	−0.14 (0.37)
	CRLB	(0.18)	(0.26)	(0.35)

(b) Means of 1,000 simulated estimates of ρ (with the standard error of a single estimate) using the MLE. The standard errors of a single observation are consistent with the values given in the CRLB line.

between class probabilities to be p , for $p = 0.05, 0.1$ and 0.2 . The value of ρ was kept constant at 0.1 , and the parameters β were as in the first experiment. The results of estimation using $\hat{\rho}_C$ are given in Table 2(a). In the first row, model $M1$, the design matrix was also constructed as before, with the first column taken to be $\mathbf{1}$, and with the remaining columns chosen to have independent standard normal entries. For the second row, model $M2$, the second column of the design matrix had elements $+1$ for the vertices in the first of the classes, and -1 for vertices in the second, corresponding to an inter-block contrast, which significantly changes the bias correction $\text{Tr}\{H_\rho C\}$. In each cell, the results of 1,000 replicates are given. At this accuracy, a small negative bias in the estimates of ρ using $\hat{\rho}_C$ can be discerned, although it is negligible when compared to the standard errors of the estimates. The approximate ideal standard errors derived from (3.7) for $\rho = 0.1$ are given in the final row. In Table 2(b), the results of using maximum likelihood are given for comparison; here, the bias is clearly visible.

Table 3 gives the results of assessing the variability of $\hat{\rho}_W$ using the permutation algorithm based on (5.12)–(5.14). The experiments were carried out in the context of the network disturbance model $M2$ based on an Erdős–Rényi mixture graph. In order to make a sensible comparison with the variability of the estimates $\hat{\rho}_W$ derived from repeatedly simulating data from the model, it was necessary to apply the permutation algorithm to a realization of the model for which the original estimate $\hat{\rho}_W$ was reasonably close to the true value $\rho_0 = 0.1$. Because of the substantial variability in $\hat{\rho}_W$, this was not the case in most realizations. In each experiment, the simulated variability of $\hat{\rho}_W^\pi$ agreed well with the variability of $\hat{\rho}_W$ under the model, and the value of $\hat{\psi}_{\rho_0}^W$ was also quite similar.

In Figure 1, a histogram of 1,000 values of $\hat{\rho}_W$, simulated from the model $M2$, but now with $\rho = 0.3$, is shown next to a histogram of 1,000 values of $\hat{\rho}_W^\pi$, generated using the permutation method described above from a realization of the model $M2$ in which $\hat{\rho}_W = 0.302$. The two histograms are very similar, and neither is substantially skewed, even though the value 0.3 chosen for ρ is close enough to the upper limit of 1 for some compression in the upper tail to be expected.

TABLE 3 Means $E(\hat{\rho}_W^\pi)$ and standard errors $SE(\hat{\rho}_W^\pi)$ of 1,000 simulated estimates of $\hat{\rho}_W$, using the permutation procedure, for the model $M2$ with $\rho = 0.1$, together with the scale of variation $\hat{\psi}_{\rho_0}^W$, calculated using (5.11), (A.25) and (A.30), and the corresponding standard errors $SE(\hat{\rho}_W)$ of $\hat{\rho}_W$ from Table 2.

p	$\hat{\rho}_W$	$E(\hat{\rho}_W^\pi)$	$SE(\hat{\rho}_W^\pi)$	$\hat{\psi}_{\rho_0}^W$	$SE(\hat{\rho}_W)$
0.05	0.094	0.092	0.19	0.17	0.21
0.1	0.107	0.112	0.29	0.26	0.29
0.2	0.091	0.073	0.43	0.42	0.42

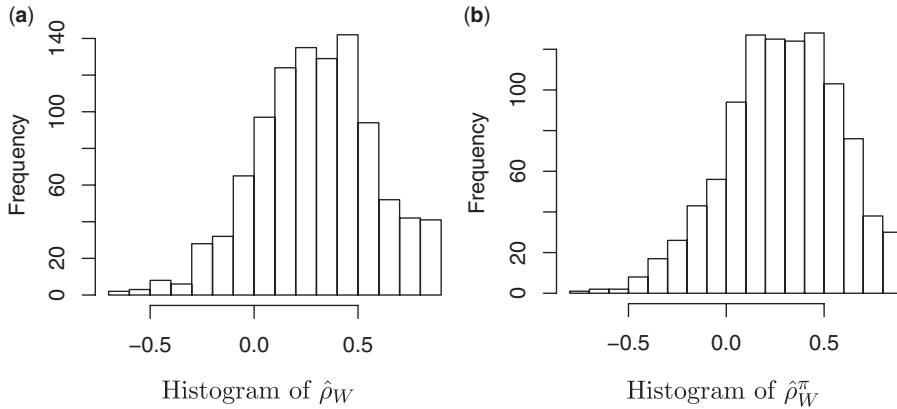


FIG. 1. Histogram of 1,000 values of $\hat{\rho}_W$ for the model $M2$ with $\rho = 0.3$, together with a histogram of 1,000 values of the permutation estimates $\hat{\rho}_W^\pi$ derived from a realization of the same model in which $\hat{\rho}_W = 0.302$

6. Network effects models

The network effects model (1.4) can be written in the form

$$y = K_\rho^{-1}X\beta + \varepsilon, \quad (6.1)$$

where the errors ε in y have the same dependence as in the network disturbance model, but the *model structure* explaining y also changes with ρ . The corresponding log-likelihood, as in [1, Equation (1.7)], is given by

$$l(\theta; y) := -\frac{n}{2}\{\log(\sigma^2) + \log(2\pi)\} - \frac{1}{2\sigma^2}(K_\rho y - X\beta)^\top (K_\rho y - X\beta) + \log \det K_\rho, \quad (6.2)$$

where $\theta := (\beta, \sigma^2, \rho)$. Since the structural element also contains information about ρ , this is reflected in the estimation procedures. In particular, the estimate of β using maximum likelihood, for a given value of ρ , is

$$\hat{\beta}_\rho(y) := (X^\top X)^{-1}XK_\rho y,$$

as can be deduced from ordinary least squares (2.1) applied to the data $y' := K_\rho y$. As a result, if the true model is $M_{\rho'}$, for $\rho' \neq \rho$, we have

$$\mathbb{E}_{\rho'} \hat{\beta}_\rho(\mathbf{y}) = (X^\top X)^{-1} X K_\rho K_{\rho'}^{-1} X \beta. \quad (6.3)$$

This expectation is in general different from β , in which case the MLE of β is likely to be biased if ρ is incorrectly estimated. Hence, in order to get good estimates of the structural component β , it is important to have an accurate estimate of ρ .

Now, differentiating (6.2) with respect to ρ , the ρ -component of the score function analogous to (3.1) takes the form

$$\tilde{s}^{(\rho)}(\theta; y) = \sigma^{-2} (K_\rho y - X\beta)^\top Z_\rho K_\rho y - \text{Tr}\{Z_\rho\}, \quad (6.4)$$

where $Z_\rho := WK_\rho^{-1}$. Its variance at θ_0 under the model with $\theta = \theta_0$ is now given by

$$V_1 := \text{Var}_{\theta_0} \tilde{s}^{(\rho)}(\theta_0; \mathbf{y}) = \text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0} Z_{\rho_0}^\top\} + \sigma_0^{-2} |Z_{\rho_0} X \beta_0|^2, \quad (6.5)$$

containing an extra term, as compared to (3.6), that involves the parameter β_0/σ_0 . This implies that the lower bound $1/V_1$ for the variance of an unbiased estimator is smaller than in the analogous network disturbance model. This does not, of itself, mean that estimation of ρ can be made more precise, since there is no guarantee that the lower bound can be attained; instead, a closer analysis is needed.

As an example, suppose that X consists solely of the column $\mathbf{1}$, with parameter β , and that the weight matrix W satisfies $W\mathbf{1} = \mathbf{1}$. Then the additional term is just $n(\beta_0/(1 - \rho_0)\sigma_0)^2$, a quantity that grows with n . On the other hand, the network effects model (6.1) translates into

$$y = (1 - \rho)^{-1} \beta \mathbf{1} + \varepsilon = \tilde{\beta} \mathbf{1} + \varepsilon,$$

which can also be viewed as a network disturbance model, with a new structural parameter $\tilde{\beta} = \beta/(1 - \rho)$ in place of β . Thus $\tilde{\beta}$ can be accurately estimated, if the network is large enough, as can be deduced from (2.11), but, as before, any estimator of ρ exhibits substantial variability if the network is dense. This explains the simulation results observed in [5], in which the maximum likelihood estimators of both ρ and the coefficient of the constant exhibited substantial variability, and were strongly negatively correlated. In the same experiments, the remaining $m - 1$ columns X' of X had independent, normally distributed entries with mean zero and variance 1. In this case, with $\beta_1/(1 - \rho)$ replaced by a parameter $\tilde{\beta}$ and with $\beta' := (\beta_2, \dots, \beta_m)^\top$, the extra term in the variance of the score function reduces to

$$\sigma^{-2} \mathbb{E}\{|Z_\rho X' \beta'|^2\} = (|\beta'|/\sigma)^2 \text{Tr}\{Z_\rho Z_\rho^\top\},$$

(where the expectation is with respect to the random choice of X'), and, as in (3.12), this does not become large with n for uniformly dense graphs. Hence, there are still substantial limitations on the precision of estimation of ρ , for Cramér–Rao reasons, and, in consequence, on the estimation of $\beta_1 = \tilde{\beta}(1 - \rho)$ also.

The same phenomenon makes itself felt if the model structure X is more complicated, but the columns of WX all belong to the linear span of the columns of X ; $WX = X\Gamma$, for an $m \times m$ matrix Γ . In this case,

the model structure $K_\rho^{-1}X\beta$ can be expressed as

$$\begin{aligned} K_\rho^{-1}X\beta &= (I - \rho W)^{-1}X\beta = X\beta + \sum_{l \geq 1} \rho^l W^l X\beta \\ &= X\beta + \sum_{l \geq 1} \rho^l X \Gamma^l \beta = X(I - \rho \Gamma)^{-1} \beta =: X\gamma, \end{aligned} \quad (6.6)$$

where the new parameter $\gamma := (I - \rho \Gamma)^{-1} \beta$. Once again, this yields a network disturbance model, in which the parameter γ can be accurately estimated if n is large, whereas, in dense networks, the estimator of ρ may exhibit substantial variability. For the original parameter β , this would express itself in variable, but strongly dependent, estimators of the components of $\beta = (I - \rho \Gamma)\gamma$, which are linked through the common value of the estimator of ρ .

In general, WX can be written in the form

$$WX = X\Gamma + E, \quad (6.7)$$

where E is an $n \times m$ matrix, whose columns are orthogonal to those of X , which measures how far WX lies from the linear span of the columns of X . In practice, except for the model in which only the overall mean is to be estimated, and $W\mathbf{1} = \mathbf{1}$ by construction, it is usual to expect that E is not precisely the zero matrix. It is shown in Appendix A.6 that, if the quantity

$$\lambda_{\max}\{E^\top E\} + \text{Tr}\{W^2 + WW^\top\} = \lambda_{\max}\{(WX)^\top(I - H)WX\} + \text{Tr}\{W^2 + WW^\top\}$$

is not large, then neither ρ nor β can be accurately estimated.

Conclusion 6: In the network effects model, if the columns of WX (almost) belong to the space generated by the columns of X , then estimating ρ and the parameters β in uniformly dense networks may both be problematic.

7. Conclusion

Maximum likelihood estimation of the correlation parameter ρ in the network disturbance and network effects models suffers from substantial bias and variability when the underlying graph is dense, as has been documented, for instance, in [3–7]. In the context of the network disturbance model (2.7), we have shown that, in dense networks, there is a lower bound to the variability of *any* reasonable estimator of ρ , as a consequence of Cramér–Rao theory, and that increasing the size of the network does not reduce the variability, if the density remains constant. We also demonstrate that, under similar circumstances, the maximum likelihood estimator is significantly biased. For certain ill conditioned networks, these effects are particularly marked. In Section 5.1, we propose a simple estimator of ρ that has little bias, though its variability in dense networks is still substantial, because all estimators are subject to the limitations imposed by the Cramér–Rao lower bound. We also suggest a permutation algorithm for assessing the variability of the estimator, on the basis of data consisting of a single network. The effectiveness of the new estimator is illustrated by simulation in a few simple settings in Section 5.2. There is a discussion of network effects models in Section 6. Here, in circumstances in which ρ cannot be accurately estimated, the estimates of the structural parameters β are also affected. An explicit quantity is given that indicates

whether such problems are likely to occur. It is expressed in terms of the design matrix X and the weight matrix W .

Acknowledgements

A.D.B. thanks the mathematics departments of the University of Melbourne and Monash University, for their kind hospitality while part of the work was undertaken. Both authors express their warm appreciation of the many helpful suggestions by the referees; these have greatly improved the presentation.

Funding

GR was supported in part by Engineering and Physical Sciences Research Council grants EP/T018445/1 and EP/R018472/1.

REFERENCES

1. DOREIAN, P. (1989) Network autocorrelation models: problems and prospects. *Spatial Statistics: Past, Present, Future* (D. A. Griffin ed.). Ann Arbor: Michigan Document Services, pp. 369–389.
2. ORD, K. (1975) Estimation methods for models of spatial interaction. *J. Am. Stat. Assoc.*, **70**, 120–126.
3. FARBER, S., PÁEZ, A. & VOLZ, E. (2010) Topology, dependency tests and estimation bias in network autoregressive models. *Progress in Spatial Analysis* (A. Páez, J. Gallo, R. N. Bullung & S. Dall’erba eds). Berlin: Springer, pp. 29–57.
4. LA ROCCA, M., PORZIO, G. C., VITALE, M. P. & DOREIAN, P. (2018) Finite sample behavior of MLE in network autocorrelation models. *Classification, (Big) Data Analysis and Statistical Learning* (F. Mola, C. Conversano & M. Vichi eds). Cham: Springer, pp. 43–50.
5. MIZRUCHI, M. S. & NEUMAN, E. J. (2008) The effect of density on the level of bias in the network autocorrelation model. *Soc. Netw.*, **30**, 190–200.
6. NEUMAN, E. J. & MIZRUCHI, M. S. (2010) Structure and bias in the network autocorrelation model. *Soc. Netw.*, **32**, 290–300.
7. SMITH, T. E. (2009) Estimation bias in spatial models with strongly connected weight matrices. *Geogr. Anal.*, **41**, 307–332.
8. BAO, Y. (2013) Finite sample bias of the QMLE in spatial autoregressive models. *Econ. Theory*, **29**, 68–89.
9. YANG, Z. (2015) A general method for third-order bias and variance corrections on a nonlinear estimator. *J. Econ.*, **186**, 178–200.
10. YU, D., BAI, P. & DING, C. (2015) Adjusted quasi-maximum likelihood estimator for mixed regressive, spatial autoregressive model and its small sample bias. *Comput. Stat. Data Anal.*, **87**, 116–135.
11. KELEJIAN, H. H. & PRUCHA, I. R. (1999) A generalized moments estimator for the autoregressive parameter in a spatial model. *Int. Econ. Rev.*, **40**, 509–533.
12. DITTRICH, D., LEENDERS, R. T. A. J. & MULDER, J. (2017) Bayesian estimation of the network autocorrelation model. *Soc. Netw.*, **48**, 213–236.
13. STRUMANN, C. (2019) Hodges–Lehmann estimation of static panel models with spatially correlated disturbances. *Comput. Econ.*, **53**, 141–168.
14. RAO, C. R. (1973) *Linear Statistical Inference and its Applications*, 2nd edn. New York: Wiley.
15. HORN, R. A. & JOHNSON, C. R. (2012) *Matrix Analysis*. Cambridge, UK: Cambridge University Press.
16. SEARLE, S. R. & GRUBER, M. H. J. (1971) *Linear Models*, vol 10. New York, London: Wiley Online Library.
17. HÄRDLE, W. K. & SIMAR, L. (2015) *Applied Multivariate Statistical Analysis*, 4th edn. Heidelberg: Springer.
18. LEHMANN, E. L. & CASELLA, G. (2006) *Theory of Point Estimation*, 2nd edn. New York: Springer.

Appendix

A.1 Rewriting the likelihood equation for ρ

To minimize the quantity $F(\rho; y)$ of (2.14), it is usual to look for values of ρ such that $\frac{d}{d\rho}F(\rho; y) = 0$. First, as for (3.1),

$$\frac{d}{d\rho} \log \det K_\rho = -\text{Tr}\{Z_\rho\}. \quad (\text{A.1})$$

Then the derivative of $\hat{\sigma}_\rho^2$ in (2.8) is given by

$$n \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} = y^\top \left\{ -W^\top (I - H_\rho) K_\rho - K_\rho^\top \left(\frac{dH_\rho}{d\rho} \right) K_\rho - K_\rho^\top (I - H_\rho) W \right\} y. \quad (\text{A.2})$$

To compute $\frac{d}{d\rho}H_\rho$, note that, for any family of invertible matrices $(Y_x, x \in \mathbb{R})$ with differentiable components, $0 = \frac{d}{dx}\{Y_x^{-1}Y_x\}$, implying that

$$\frac{d}{dx}\{Y_x^{-1}\} = -Y_x^{-1} \frac{d}{dx}\{Y_x\} Y_x^{-1}. \quad (\text{A.3})$$

Hence, writing $Y_\rho := ((K_\rho X)^\top K_\rho X)$, and using $\frac{dK_\rho}{d\rho} = -W$, we have

$$\frac{dY_\rho^{-1}}{d\rho} = Y_\rho^{-1} \{X^\top W^\top K_\rho X + X^\top K_\rho^\top W X\} Y_\rho^{-1}.$$

Thus, noting for the last line that $W = Z_\rho K_\rho$ and that $K_\rho X Y_\rho^{-1} (K_\rho X)^\top = H_\rho$, we deduce that

$$\begin{aligned} \frac{dH_\rho}{d\rho} &= \frac{d}{d\rho} \left\{ K_\rho X Y_\rho^{-1} (K_\rho X)^\top \right\} \\ &= -W X Y_\rho^{-1} (K_\rho X)^\top - K_\rho X Y_\rho^{-1} (W X)^\top \\ &\quad + K_\rho X Y_\rho^{-1} \{X^\top W^\top K_\rho X + X^\top K_\rho^\top W X\} Y_\rho^{-1} (K_\rho X)^\top \\ &= -(I - H_\rho) Z_\rho H_\rho - H_\rho Z_\rho^\top (I - H_\rho). \end{aligned} \quad (\text{A.4})$$

Note also that, because H_ρ is idempotent, which implies that $(I - H_\rho)H_\rho = 0$, we have

$$\text{Tr}((Z_\rho H_\rho)^\top (I - H_\rho)) = \text{Tr}((I - H_\rho) Z_\rho H_\rho) = \text{Tr}(Z_\rho H_\rho (I - H_\rho)) = 0. \quad (\text{A.5})$$

Combining (A.2) and (A.4), it follows that

$$n \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} = -(K_\rho y)^\top (I - H_\rho) (Z_\rho^\top + Z_\rho) (I - H_\rho) K_\rho y, \quad (\text{A.6})$$

and thus, writing $U(\rho; y) := \frac{d}{d\rho} F(\rho; y)$ and using (A.1),

$$U(\rho; y) = \frac{1}{\hat{\sigma}_\rho^2(y)} \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} - \frac{2}{n} \frac{d}{d\rho} \log \det K_\rho \quad (\text{A.7})$$

$$= -\frac{1}{n\hat{\sigma}_\rho^2(y)} (K_\rho y)^\top (I - H_\rho)(Z_\rho^\top + Z_\rho)(I - H_\rho)K_\rho y + \frac{2}{n} \sum_{l \geq 1} \rho^l \text{Tr}\{W^{l+1}\}, \quad (\text{A.8})$$

as given in (4.2).

A.2 The magnitude of $nU'(\rho; y)$

From (A.7) and (A.8), and since, for any $n \times n$ matrix M and n -vector y , $y^\top M y = y^\top M^\top y$, it follows that

$$\begin{aligned} U'(\rho; y) &= \frac{1}{\hat{\sigma}_\rho^2(y)} \frac{d^2 \hat{\sigma}_\rho^2(y)}{d\rho^2} - \left\{ \frac{1}{\hat{\sigma}_\rho^2(y)} \frac{d\hat{\sigma}_\rho^2(y)}{d\rho} \right\}^2 + \frac{2}{n} \sum_{l \geq 1} l \rho^{l-1} \text{Tr}\{W^{l+1}\} \\ &= -\frac{2}{n\hat{\sigma}_\rho^2(y)} \frac{d}{d\rho} \{y^\top K_\rho^\top \{(I - H_\rho)Z_\rho(I - H_\rho)\} K_\rho y\} \\ &\quad - \left\{ U(\rho; y) - \frac{2}{n} \sum_{l \geq 1} \rho^l \text{Tr}\{W^{l+1}\} \right\}^2 + \frac{2}{n} \sum_{l \geq 1} l \rho^{l-1} \text{Tr}\{W^{l+1}\}, \end{aligned} \quad (\text{A.9})$$

invoking (A.6) for the first term. Using $\frac{d}{d\rho} K_\rho = -W = -Z_\rho K_\rho$ and $\frac{d}{d\rho} Z_\rho = W \frac{d}{d\rho} \{K_\rho^{-1}\} = Z_\rho^2$, from (A.3), and again using $y^\top M y = y^\top M^\top y$, we can express the derivative of the quadratic form in (A.9) in terms of the quantities H_ρ , K_ρ , Z_ρ and $\frac{d}{d\rho} H_\rho$ as

$$\begin{aligned} &\frac{d}{d\rho} \{y^\top K_\rho^\top \{(I - H_\rho)Z_\rho(I - H_\rho)\} K_\rho y\} \\ &= y^\top K_\rho^\top \left\{ -Z_\rho^\top (I - H_\rho)(Z_\rho + Z_\rho^\top)(I - H_\rho) - \frac{dH_\rho}{d\rho} (Z_\rho + Z_\rho^\top)(I - H_\rho) \right. \\ &\quad \left. + (I - H_\rho)Z_\rho^2(I - H_\rho) \right\} K_\rho y. \end{aligned} \quad (\text{A.10})$$

The middle term in (A.9) is expressed in terms of the same quantities H_ρ , K_ρ , Z_ρ and $\frac{d}{d\rho} H_\rho$, together with W , using (A.2). The final term in (A.9), much as for (4.6), gives $2n^{-1} \text{Tr}\{Z_\rho^2\}$, and, from (A.4),

$$\begin{aligned} &y^\top K_\rho^\top \frac{dH_\rho}{d\rho} (Z_\rho + Z_\rho^\top)(I - H_\rho)K_\rho y \\ &= -y^\top K_\rho^\top \{(I - H_\rho)Z_\rho H_\rho + H_\rho Z_\rho^\top (I - H_\rho)\} (Z_\rho + Z_\rho^\top)(I - H_\rho)K_\rho y. \end{aligned} \quad (\text{A.11})$$

The resulting expressions for $U'(\rho; y)$ are not easy to understand.

It is somewhat easier to consider, as a typical value, the expectation $n\mathbb{E}_{\rho_0} U'(\rho_0; \mathbf{y})$ on the model M_{ρ_0} . Approximating $\hat{\sigma}_\rho^2(\mathbf{y})$ by σ^2 , the expectation of n times the first and third terms in (A.9) is approximated

by

$$\mathbb{E}_{\rho_0} \left\{ -\frac{2}{\sigma^2} \frac{d}{d\rho} \{ \mathbf{y}^\top K_\rho^\top \{ (I - H_\rho) Z_\rho (I - H_\rho) \} K_\rho \mathbf{y} \} \Big|_{\rho=\rho_0} \right\} + 2\text{Tr}\{Z_{\rho_0}^2\}.$$

This expression, using (3.3) together with (A.10) and (A.11) and properties of the trace, and also observing that $(I - H_\rho)K_\rho X = 0$, gives

$$\begin{aligned} & 2\text{Tr}\{Z_{\rho_0}^\top (I - H_{\rho_0})(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0}) \\ & \quad - \{(I - H_{\rho_0})Z_{\rho_0}H_{\rho_0} + H_{\rho_0}Z_{\rho_0}^\top(I - H_{\rho_0})\}(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0}) \\ & \quad - (I - H_{\rho_0})Z_{\rho_0}^2(I - H_{\rho_0}) + Z_{\rho_0}^2\} \\ & = \text{Tr}\{(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})\} \\ & \quad + 2\text{Tr}\{-Z_{\rho_0}H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0}) + H_{\rho_0}Z_{\rho_0}^2\} =: \Delta_0, \end{aligned} \quad (\text{A.12})$$

with the first trace equal to $\frac{1}{2}\tau_0^2$, as given in (4.8). The middle term in (A.9), multiplied by n , gives

$$n^{-1}\text{Var}_{\rho_0}(nU(\rho; \mathbf{y})) + n^{-1}(n\mathbb{E}_{\rho_0}U(\rho; \mathbf{y}) - 2\text{Tr}\{Z_0\})^2 \approx n^{-1}(\tau_0^2 + 4\text{Tr}\{(I - H_{\rho_0})Z_{\rho_0}\}^2),$$

using (4.7) and (4.8). In the circumstances discussed in Section 4.3, Δ_0 and τ_0^2 are seen to be typically comparable to $\text{Tr}\{Z_{\rho_0}Z_{\rho_0}^\top\}$, so that this last contribution is of relative order $O(n^{-1})$ as compared to Δ_0 . Hence, in such circumstances, the expectation $n\mathbb{E}_{\rho_0}U'(\rho_0; \mathbf{y})$ on the model M_{ρ_0} is approximately given by Δ_0 .

A.3 Details of the MLE when structure is present

Suppose that X consists solely of the n -vector $\mathbf{1}$, having all elements equal to 1, and that there are no isolated vertices. Suppose also that W is chosen to have $\sum_{j=1}^n W_{ij} = 1$ for all j , so that $W\mathbf{1} = \mathbf{1}$. Then $H_\rho = H_* := n^{-1}\mathbf{1}\mathbf{1}^\top$ for all ρ , because $K_\rho X = (I - \rho W)\mathbf{1} = (1 - \rho)\mathbf{1}$, and hence $WH_\rho = H_*$ and $Z_\rho H_\rho = (1 - \rho)^{-1}H_*$. The matrix product in (4.8) can be multiplied out to give

$$\begin{aligned} \tau_0^2 &= 2\text{Tr}\{(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})(Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_{\rho_0})\} \\ &= 2\text{Tr}\{2(Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top) - 2H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)^2 + [H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)]^2\} \\ &= 4\text{Tr}\{(Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top)(I - H_{\rho_0})\} - 2\text{Tr}\{H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)^2 - [H_{\rho_0}(Z_{\rho_0} + Z_{\rho_0}^\top)]^2\}. \end{aligned} \quad (\text{A.13})$$

With the particular choices of X and W , and repeatedly using the properties of the trace, this gives

$$\tau_0^2 = 4\text{Tr}\{(Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top)(I - H_*)\} + 2\left\{\frac{1}{(1 - \rho_0)^2} - \frac{1}{n}|Z_{\rho_0}^\top \mathbf{1}|^2\right\},$$

where, for instance, the final term is just $-\text{Tr}\{H_*Z_{\rho_0}Z_{\rho_0}^\top\}$, whereas, for example,

$$\text{Tr}\{H_*Z_{\rho_0}^\top Z_{\rho_0}\} = \frac{1}{1 - \rho_0}\text{Tr}\{H_*Z_{\rho_0}\} = \frac{1}{1 - \rho_0}\text{Tr}\{Z_{\rho_0}H_*\} = \frac{1}{(1 - \rho_0)^2}.$$

In much the same way, and in particular because $Z_\rho H_* = (1 - \rho)^{-1} H_*$ and $H_*(I - H_*) = 0$, we obtain

$$2\text{Tr}\{-Z_{\rho_0} H_* (Z_{\rho_0} + Z_{\rho_0}^\top)(I - H_*) + H_* Z_{\rho_0}^2\} = 2\text{Tr}\{Z_{\rho_0}^2 H_*\} = \frac{2}{(1 - \rho_0)^2},$$

giving the formula for Δ_0 in (4.15).

Note that

$$\frac{1}{n} |Z_{\rho_0}^\top \mathbf{1}|^2 \geq \left\{ \frac{1}{n} \mathbf{1}^\top Z_{\rho_0}^\top \mathbf{1} \right\}^2 = \frac{1}{(1 - \rho_0)^2},$$

by Cauchy–Schwarz, giving the upper bound for τ_0^2 in (4.15). Note also that, again using $Z_\rho H_\rho = (1 - \rho)^{-1} H_*$, and because $\text{Tr}\{H_*\} = 1$ and $H_*^2 = H_*$, it follows that

$$\begin{aligned} \frac{1}{4} \tau_0^2 &= \text{Tr}\{Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2\} - \text{Tr}\{H_* Z_{\rho_0} (H_* Z_{\rho_0})^\top\} - \text{Tr}\{Z_{\rho_0}^2 H_*\} \\ &\leq \text{Tr}\{Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2\} - \frac{1}{(1 - \rho_0)^2}, \end{aligned} \quad (\text{A.14})$$

so that τ_0^2 is indeed smaller than $4\text{Tr}\{Z_{\rho_0} Z_{\rho_0}^\top + Z_{\rho_0}^2\}$.

Suppose now that there is more structure in the regression, so that X consists of columns in addition to the first column $\mathbf{1}$. If we still assume that $W\mathbf{1} = \mathbf{1}$, the basic message remains the same. Because the matrix H_ρ is idempotent, its spectral decomposition can be written as

$$H_\rho = H_* + \sum_{l=2}^m x_\rho^{(l)} (x_\rho^{(l)})^\top, \quad (\text{A.15})$$

where $H_* := n^{-1} \mathbf{1} \mathbf{1}^\top$ is as before, and $x_\rho^{(2)}, \dots, x_\rho^{(m)}$ are *orthonormal* vectors orthogonal to $\mathbf{1}$ that, with $\mathbf{1}$, span the space generated by the columns of $K_\rho X$. If, for instance, it is assumed that each of the vectors $x_\rho^{(l)}$, $2 \leq l \leq m$, has components bounded in modulus by $cn^{-1/2}$, for some fixed $c \geq 1$ —as is the case for $x^{(1)} = n^{-1/2} \mathbf{1}$ —then the traces of the matrices $Z_{\rho_0}^2 H_{\rho_0}$, $Z_{\rho_0}^\top Z_{\rho_0} H_{\rho_0}$, $(Z_{\rho_0} H_{\rho_0})^2$ and $Z_{\rho_0}^\top H_{\rho_0} Z_{\rho_0} H_{\rho_0}$ can all be shown to be bounded by $m^2 c^4 / (1 - |\rho|)^2$, which is not large if m and c are moderate and $|\rho|$ is not too close to 1. The argument runs as follows.

In view of (4.16), and writing $x^{(1)} := n^{-1/2} \mathbf{1}$, the trace $\text{Tr}\{MH_\rho\}$, for any $n \times n$ matrix M , can be written in the form

$$\text{Tr}\{MH_\rho\} = \sum_{l=1}^m \text{Tr}\{M x_\rho^{(l)} (x_\rho^{(l)})^\top\} = \sum_{l=1}^m (x_\rho^{(l)})^\top M x_\rho^{(l)}. \quad (\text{A.16})$$

Thus, using (2.6), it follows that, for any $1 \leq l, l' \leq m$,

$$\begin{aligned} |Z_\rho x_\rho^{(l)}| &= \left| \sum_{r \geq 0} W^{r+1} \rho^r x_\rho^{(l)} \right| \\ &\leq \sum_{r \geq 0} |\rho|^r c n^{-1/2} |W^{r+1} \mathbf{1}| = \frac{c}{n^{1/2}(1 - |\rho|)} |\mathbf{1}| = \frac{c}{1 - |\rho|}; \end{aligned} \quad (\text{A.17})$$

that

$$\begin{aligned}
 |(x_\rho^{(l)})^\top Z_\rho x_\rho^{(l')}| &= \left| \sum_{r \geq 0} (x_\rho^{(l)})^\top W^{r+1} \rho^r x_\rho^{(l')} \right| \\
 &\leq \left| \sum_{r \geq 0} |\rho|^r (cn^{-1/2} \mathbf{1}^\top) W^{r+1} (cn^{-1/2} \mathbf{1}) \right| = c^2 n^{-1} \mathbf{1}^\top \mathbf{1} \sum_{r \geq 0} |\rho|^r \\
 &= \frac{c^2}{1 - |\rho|};
 \end{aligned} \tag{A.18}$$

and that

$$\begin{aligned}
 |(x_\rho^{(l)})^\top Z_\rho^2 x_\rho^{(l')}| &= |(x_\rho^{(l)})^\top W^2 (1 - \rho W)^{-2} x_\rho^{(l')}| = \left| \sum_{r \geq 0} (x_\rho^{(l)})^\top (r+1) W^{r+2} \rho^r x_\rho^{(l')} \right| \\
 &\leq c^2 n^{-1} \mathbf{1}^\top \mathbf{1} \sum_{r \geq 0} (r+1) |\rho|^r = \frac{c^2}{(1 - |\rho|)^2}.
 \end{aligned} \tag{A.19}$$

Then, for instance, from (A.16) and (A.19), we have

$$\text{Tr}\{Z_{\rho_0}^2 H_{\rho_0}\} = \sum_{l=1}^m (x_\rho^{(l)})^\top Z_{\rho_0}^2 x_\rho^{(l)} \leq \frac{mc^2}{(1 - |\rho|)^2};$$

from (A.16) and (A.17), we have

$$\text{Tr}\{Z_{\rho_0}^\top Z_{\rho_0} H_{\rho_0}\} = \sum_{l=1}^m (x_\rho^{(l)})^\top Z_{\rho_0}^\top Z_{\rho_0} x_\rho^{(l)} = \sum_{l=1}^m |Z_{\rho_0} x_\rho^{(l)}|^2 \leq \frac{mc^2}{(1 - |\rho|)^2};$$

and, from (A.16) and (A.18), we have

$$\begin{aligned}
 \text{Tr}\{Z_{\rho_0}^\top H_{\rho_0} Z_{\rho_0} H_{\rho_0}\} &= \sum_{l=1}^m (x_\rho^{(l)})^\top Z_{\rho_0}^\top H_{\rho_0} Z_{\rho_0} x_\rho^{(l)} = \sum_{l=1}^m \text{Tr}\{Z_{\rho_0} x_\rho^{(l)} (x_\rho^{(l)})^\top Z_{\rho_0}^\top H_{\rho_0}\} \\
 &= \sum_{l=1}^m \sum_{l'=1}^m (x_\rho^{(l')})^\top Z_{\rho_0} x_\rho^{(l)} (x_\rho^{(l)})^\top Z_{\rho_0}^\top x_\rho^{(l')} \leq \frac{m^2 c^4}{(1 - |\rho|)^2}.
 \end{aligned}$$

By similar calculations, each of the traces listed above is bounded in modulus by the quantity $m^2 c^4 / (1 - |\rho|)^2$.

A.4 Estimation of β

From (2.8) and (2.11), the distribution of $\hat{\beta}_\rho$ on the model M_ρ has covariance matrix $\sigma^2 (X^\top S_\rho X)^{-1}$, where $S_\rho := K_\rho^\top K_\rho$. The variance of a linear combination $a^\top \hat{\beta}_\rho$, for a given m -dimensional unit vector a , is thus

$$\sigma^2 a^\top (X^\top S_\rho X)^{-1} a \leq \sigma^2 \lambda_{\max}((X^\top S_\rho X)^{-1}) = \sigma^2 / \lambda_{\min}(X^\top S_\rho X), \tag{A.20}$$

where $\lambda_{\min}(M) \leq \lambda_{\max}(M)$ denote the smallest and largest eigenvalues of M , both positive if M is positive definite symmetric. Letting \tilde{u}_X denote the unit n -vector $Xu/|Xu|$, and using Rayleigh–Ritz twice, we have

$$\begin{aligned}\lambda_{\min}(X^\top S_\rho X) &= \min_{u \in \mathbb{R}^m: |u|=1} u^\top X^\top S_\rho X u = \min_{u \in \mathbb{R}^m: |u|=1} |Xu|^2 \tilde{u}_X^\top S_\rho \tilde{u}_X \\ &\geq \lambda_{\min}(S_\rho) \min_{u \in \mathbb{R}^m: |u|=1} |Xu|^2 = \lambda_{\min}(X^\top X) \lambda_{\min}(S_\rho).\end{aligned}\quad (\text{A.21})$$

Hence, for fixed σ^2 , provided that $\lambda_{\min}(S_\rho)$ is not close to zero, (A.20) and (A.21) imply that accurate estimation of all linear combinations of β depends on having a large value of $\lambda_{\min}(X^\top X)$, which is the usual condition for ordinary least squares. In particular, if the network size n becomes large, but $\lambda_{\min}(S_\rho)$ remains uniformly bounded away from zero, the estimate of β is consistent if $\lambda_{\min}(X^\top X) \rightarrow \infty$. Example 1 in Section 4.4 shows that the latter condition alone is not enough to guarantee consistent estimation of β .

If the true value of ρ is ρ_0 , the distribution of $\hat{\beta}_\rho(\mathbf{y})$, as determined by substituting $\mathbf{y} = X\beta + K_{\rho_0}^{-1}\mathbf{v}$ for \mathbf{y} into (2.8), is multivariate normal, with mean β and covariance matrix $\sigma^2 \Sigma(\rho, \rho_0)$, where

$$\begin{aligned}\Sigma(\rho, \rho_0) &:= (X^\top S_\rho X)^{-1} (K_\rho X)^\top K_\rho K_{\rho_0}^{-1} (K_{\rho_0}^{-1})^\top K_\rho^\top K_\rho X (X^\top S_\rho X)^{-1} \\ &= (X^\top S_\rho X)^{-1} X^\top S_\rho (S_{\rho_0})^{-1} S_\rho X (X^\top S_\rho X)^{-1}.\end{aligned}\quad (\text{A.22})$$

Writing $\tilde{X}_\rho := S_\rho^{1/2} X (X^\top S_\rho X)^{-1}$, where $M^{1/2}$ denotes the non-negative square root of a non-negative definite matrix M , and setting $S(\rho, \rho_0) := S_\rho^{1/2} (S_{\rho_0})^{-1} S_\rho^{1/2}$, (A.22) gives

$$\Sigma(\rho, \rho_0) = \tilde{X}_\rho^\top S(\rho, \rho_0) \tilde{X}_\rho. \quad (\text{A.23})$$

Using Rayleigh–Ritz, for any unit m -vector a , we thus have

$$\begin{aligned}a^\top \Sigma(\rho, \rho_0) a &= (\tilde{X}_\rho a)^\top S(\rho, \rho_0) \tilde{X}_\rho a \leq |\tilde{X}_\rho a|^2 \lambda_{\max}(S(\rho, \rho_0)) \\ &= \{a^\top (X^\top S_\rho X)^{-1} a\} \lambda_{\max}(S(\rho, \rho_0)).\end{aligned}$$

Thus, if the true model is M_{ρ_0} , estimating a linear combination $a^\top \beta$ using $\hat{\beta}_\rho$ instead of $\hat{\beta}_{\rho_0}$ gives an estimator $a^\top \hat{\beta}_\rho$ whose variance is at most $\lambda_{\max}(S(\rho, \rho_0))$ times the value, given in (A.20), that it would have had, were M_ρ the true model. Hence, if estimation on the model M_ρ , using $\hat{\beta}_\rho$, is accurate for any ρ , then estimating using $\hat{\beta}_\rho$ on the model M_{ρ_0} is also accurate, provided that $\lambda_{\max}(S(\rho, \rho_0))$ is not large.

Note that, if $n(W) := \sqrt{\lambda_{\max}(W^\top W)}$ denotes the operator norm of the matrix W , then, by Rayleigh–Ritz, because $|Wy| \leq n(W)|y|$ for all y , the largest eigenvalue of $S(\rho, 0) = I - \rho(W + W^\top) + \rho^2 W^\top W$ is at least $(|\rho|n(W) - 1)^2$, and this is large, for moderate non-zero values of ρ , if $n(W)$ is large. Hence, if $\lambda_{\max}(W^\top W)$ is large, this indicates that there may possibly be problems with estimating β .

A.5 The scale of precision of estimation using T_ρ^C

As discussed in Section 5, a plausible measure of the scale of precision of estimation using T_ρ^C is given by $\hat{\psi}_{\rho_0}^C := \hat{\tau}_{\rho_0}^C / |\hat{\Delta}_{\rho_0}^C|$, to be estimated by $\hat{\tau}_{\hat{\rho}_C}^C / |\hat{\Delta}_{\hat{\rho}_C}^C|$, where

$$(\hat{\tau}_{\rho_0}^C)^2 := \sigma^{-4} \text{Var}_{\rho_0}(U_{\rho_0}^C(\mathbf{y})) \quad \text{and} \quad \hat{\Delta}_{\rho_0}^C := -\sigma^{-2} \mathbb{E}_{\rho_0} \left\{ \frac{dU_\rho^C(\mathbf{y})}{d\rho} \Big|_{\rho=\rho_0} \right\},$$

and U^C is as defined in (5.8).

To find $(\hat{\tau}_{\rho_0}^C)^2$, note that, from (5.4) and (2.8), since $\mathbf{y} = X\beta + K_{\rho_0}^{-1}\mathbf{v}$ under M_{ρ_0} , we have

$$U_{\rho_0}^C(\mathbf{y}) = \mathbf{v}^\top (Q_C(\rho_0, \rho_0) + n^{-1}(I - H_{\rho_0})\text{Tr}\{H_{\rho_0}C\})\mathbf{v} = \mathbf{v}^\top \widehat{Q}(C, \rho_0)\mathbf{v},$$

where $Q_C(\rho, \rho_0)$ is defined in (5.5), and

$$\widehat{Q}(C, \rho_0) := (I - H_{\rho_0})C(I - H_{\rho_0}) + n^{-1}(I - H_{\rho_0})\text{Tr}\{H_{\rho_0}C\}. \quad (\text{A.24})$$

Hence, replacing C by $\frac{1}{2}(C + C^\top)$, which does not change T^C , and using (3.5), it follows that

$$(\hat{\tau}_{\rho_0}^C)^2 = \sigma^{-4} \text{Var}_{\rho_0}(U_{\rho_0}^C(\mathbf{y})) = \frac{1}{2} \text{Tr}\{\{\widehat{Q}(C + C^\top, \rho_0)\}^2\}. \quad (\text{A.25})$$

In view of the factor n^{-1} in the second element of $\widehat{Q}(C, \rho_0)$, the dominant contribution is that from

$$\frac{1}{2} \text{Tr}\{\{Q_C(\rho_0, \rho_0)\}^2\} = \text{Tr}\{(I - H_{\rho_0})C(I - H_{\rho_0})(C + C^\top)\},$$

once more of the form $\text{Tr}\{Z^2 + ZZ^\top\}$, now with $Z := (I - H_{\rho_0})C(I - H_{\rho_0})$.

To find $\widehat{\Delta}_{\rho_0}^C$, observe that, since $\frac{d}{d\rho}K_\rho = -W = -Z_\rho K_\rho$ and $\frac{d}{d\rho}H_\rho$ is as in (A.4), we have

$$\begin{aligned} \frac{dT_\rho^C(\mathbf{y})}{d\rho} &= \frac{d}{d\rho}\{(K_\rho\mathbf{y})^\top (I - H_\rho)C(I - H_\rho)K_\rho\mathbf{y}\} \\ &= -(K_\rho\mathbf{y})^\top \left(Z_\rho^\top (I - H_\rho) + \frac{dH_\rho}{d\rho} \right) (C + C^\top)(I - H_\rho)K_\rho\mathbf{y} \\ &= -(X\beta + K_{\rho_0}^{-1}\mathbf{v})^\top K_\rho^\top \{ Z_\rho^\top (I - H_\rho) - (I - H_\rho)Z_\rho H_\rho \\ &\quad - H_\rho Z_\rho^\top (I - H_\rho) \} (C + C^\top)(I - H_\rho)K_\rho K_{\rho_0}^{-1}\mathbf{v}, \end{aligned} \quad (\text{A.26})$$

where $X\beta$ is not present at the second appearance of \mathbf{v} because $(I - H_\rho)K_\rho X = 0$. From (A.26) and (3.3), and because $H_\rho(I - H_\rho) = 0$, it thus follows that

$$-\mathbb{E}_{\rho_0} \left\{ \frac{dT_\rho^C(\mathbf{y})}{d\rho} \Big|_{\rho=\rho_0} \right\} = \sigma^2 \text{Tr}\{(I - H_{\rho_0})[Z_{\rho_0}^\top (I - H_{\rho_0}) - Z_{\rho_0} H_{\rho_0}](C + C^\top)\} =: \sigma^2 \widehat{\Delta}_{\rho_0,1}^C, \quad (\text{A.27})$$

say, with $X\beta$ disappearing because $\mathbb{E}\mathbf{v} = 0$. Then, from (A.4),

$$-\frac{d}{d\rho} \text{Tr}\{H_\rho C\} = -\text{Tr}\left\{ C \frac{dH_\rho}{d\rho} \right\} = \text{Tr}\{C\{(I - H_\rho)Z_\rho H_\rho + H_\rho Z_\rho^\top (I - H_\rho)\}\} =: \widehat{\Delta}_{\rho_0,2}^C, \quad (\text{A.28})$$

say, and, from (A.6), and on M_{ρ_0} ,

$$-\frac{d\hat{\sigma}_{\rho_0}^2(\mathbf{y})}{d\rho} = n^{-1} \mathbf{v}^\top (I - H_{\rho_0})(Z_{\rho_0}^\top + Z_{\rho_0})(I - H_{\rho_0})\mathbf{v},$$

so that

$$-\sigma^{-2} \mathbb{E}_{\rho_0} \left\{ \frac{d\hat{\sigma}_{\rho_0}^2(\mathbf{y})}{d\rho} \right\} \text{Tr}\{H_{\rho_0}C\} = 2n^{-1} \text{Tr}\{(I - H_{\rho_0})Z_{\rho_0}(I - H_{\rho_0})\} \text{Tr}\{H_{\rho_0}C\} =: \widehat{\Delta}_{\rho_0,3}^C, \quad (\text{A.29})$$

say. Then

$$\widehat{\Delta}_{\rho_0}^C := \sum_{l=1}^3 \widehat{\Delta}_{\rho_0, l}^C, \quad (\text{A.30})$$

where $\widehat{\Delta}_{\rho_0, l}^C$, $1 \leq l \leq 3$, are given by (A.27), (A.28) and (A.29).

A.6 Details of network effects estimation

In the network effects model, except when only the overall mean is to be estimated, and $W\mathbf{1} = \mathbf{1}$ by construction, it is usual to expect that

$$WX = HWX + (I - H)WX = X\Gamma + E, \quad (\text{A.31})$$

for a non-zero $n \times m$ matrix E . Here, $H := X(X^\top X)^{-1}X^\top$ is the projection onto the space spanned by the columns of X ; HWX is the corresponding projection of the columns of WX , and can thus be written in the form $X\Gamma$; and $E := (I - H)WX$ has columns orthogonal to the columns of X . Then, as in the calculation leading to (6.6), we obtain

$$\begin{aligned} K_\rho^{-1}X\beta &= X(I - \rho\Gamma)^{-1}\beta + \rho K_\rho^{-1}E\beta \\ &= X\gamma + \rho K_\rho^{-1}E(I - \rho\Gamma)\gamma =: X_\rho\gamma, \end{aligned}$$

say, parametrized in terms of γ . The log-likelihood, in terms of γ , becomes

$$\ell(\gamma, \sigma^2, \rho; y) = -\frac{n}{2}\{\log(\sigma^2) + \log(2\pi)\} - \frac{1}{2\sigma^2}\{K_\rho(y - X_\rho\gamma)\}^\top\{K_\rho(y - X_\rho\gamma)\} + \log \det K_\rho,$$

and differentiating with respect to ρ gives the analogue of the score function in (6.4) to be

$$\begin{aligned} \tilde{s}^{(\rho)}(\gamma, \sigma^2, \rho; y) &= \frac{1}{\sigma^2}(\{K_\rho(y - X_\rho\gamma)\}^\top Z_\rho\{K_\rho(y - X_\rho\gamma)\}) \\ &\quad + \frac{1}{\sigma^2}\{K_\rho(y - X_\rho\gamma)\}^\top G_\rho\gamma - \text{Tr}\{Z_\rho\}, \end{aligned} \quad (\text{A.32})$$

where

$$\begin{aligned} G_\rho &:= K_\rho \frac{d}{d\rho} \{\rho K_\rho^{-1}E(I - \rho\Gamma)\} \\ &= E(I - \rho\Gamma) - \rho K_\rho^{-1}WE(I - \rho\Gamma) - \rho E\Gamma. \end{aligned} \quad (\text{A.33})$$

The variance of $\tilde{s}^{(\rho)}(\gamma_0, \sigma_0^2, \rho_0; y)$ on the model with parameters $(\gamma_0, \sigma_0^2, \rho_0)$ is now

$$V_2 := \text{Tr}\{Z_{\rho_0}^2 + Z_{\rho_0}Z_{\rho_0}^\top\} + \frac{1}{\sigma_0^2}|G_{\rho_0}\gamma|^2, \quad (\text{A.34})$$

differing from (6.5) in the second term. The corresponding Cramér–Rao lower bound for the variance of any unbiased estimator of ρ is thus at least as big as $1/V_2$.

Just as for the network disturbance model, the first term in (A.34) may not be large in dense networks. To get an idea of the magnitude of the second term, consider the case where $\rho_0 = 0$, so that $G_0 = E$ and $|G_0\gamma|^2 = \gamma^\top E^\top E \gamma$. Then

$$\frac{1}{\sigma_0^2} |G_0\gamma|^2 \leq \frac{|\gamma|^2}{\sigma_0^2} \lambda_{\max}(E^\top E),$$

where $\lambda_{\max}(E^\top E)$ denotes the largest eigenvalue of the positive definite symmetric $m \times m$ matrix $E^\top E$. For instance, if the elements of E , that represent to what extent the columns of WX do not belong to the linear span of the columns of X , are uniformly of magnitude at most c/\sqrt{n} , then the elements of $E^\top E$ are at most of magnitude c^2 , and thus $\lambda_{\max}(E^\top E)$ is not large if c is not large. In such circumstances, any reasonable estimator of ρ has to have substantial variability. In practice, the model decomposition (A.31) can be applied to yield the value of

$$\lambda_{\max}(E^\top E) = \lambda_{\max}((WX)^\top (I - H)WX),$$

and, if neither this nor $\text{Tr}\{W^2 + WW^\top\}$ is large, then neither ρ nor β can be accurately estimated.