



**UNSW**  
THE UNIVERSITY OF NEW SOUTH WALES

Faculty of  
Business

**School of Economics  
UNSW, Sydney 2052  
Australia**

<http://www.economics.unsw.edu.au>

**Bayesian Covariance Matrix Estimation using a Mixture of  
Decomposable Graphical Models**

**Helen Armstrong, Christopher K. Carter, Kevin K. F. Wong and  
Robert Kohn**

**School of Economics Discussion Paper: 2007/13**

The views expressed in this paper are those of the authors and do not necessarily reflect those of the School of Economic at UNSW.

ISSN 1323-8949

ISBN 978 0 7334 2449 6

# Bayesian Covariance Matrix Estimation using a Mixture of Decomposable Graphical Models

Helen Armstrong

*School of Mathematics, University of New South Wales Sydney, Australia.*

Christopher K. Carter

*CSIRO Mathematical and Information Sciences, Sydney, Australia.*

Kevin K. F. Wong

*The Institute of Statistical Mathematics, Graduate University for Advanced Studies, Tokyo, Japan.*

Robert Kohn

*Faculty of Commerce and Economics, University of New South Wales, Sydney, Australia.*

**Summary.** Estimating a covariance matrix efficiently and discovering its structure are important statistical problems with applications in many fields. This article takes a Bayesian approach to estimate the covariance matrix of Gaussian data. We use ideas from Gaussian graphical models and model selection to construct a prior for the covariance matrix that is a mixture over all decomposable graphs, where a graph means the configuration of nonzero off-diagonal elements in the inverse of the covariance matrix. Our prior for the covariance matrix is such that the probability of each graph size is specified by the user and graphs of equal size are assigned equal probability. Most previous approaches assume that all graphs are equally probable. We give empirical results that show the prior that assigns equal probability over graph sizes outperforms the prior that assigns equal probability over all graphs, both in identifying the correct decomposable graph and in more efficiently estimating the covariance matrix. The advantage is greatest when the number of observations is small relative to the dimension of the covariance matrix. The article also shows empirically that there is minimal change in statistical efficiency in using the mixture over decomposable graphs prior for estimating a general covariance compared to the Bayesian estimator by Wong et al. (2003), even when the graph of the covariance matrix is nondecomposable. However, our approach has some important advantages over that of Wong et al. (2003). Our method requires the number of decomposable graphs for each graph size. We show how to estimate these numbers using simulation and that the simulation results agree with analytic results when such results are known. We also show how to estimate the posterior distribution of the covariance matrix using Markov chain Monte Carlo with the elements of the covariance matrix integrated out and give empirical results that show the sampler is computationally efficient and converges rapidly. Finally, we note that both the prior and the simulation method to evaluate the prior apply generally to any decomposable graphical model.

**KEY WORDS:** Covariance selection; Graphical models; Reduced conditional sampling; Variable selection

## 1. Introduction

Estimating a covariance matrix efficiently is an important statistical problem with many applications, such as multivariate regression, cluster analysis, factor analysis, and discriminant analysis; see, for example, Mardia et al. (1979). Such applications are used in the fields of Business, Engineering, and the physical and social sciences. It is also of considerable interest to understand the graphical structure of the covariance matrix because it is directly interpretable in terms of the partial correlations of the underlying multivariate distribution. By the graph of the covariance matrix we mean the pattern of nonzero off diagonal elements in the inverse of the covariance matrix, also called the concentration matrix (see Lauritzen 1996, Chapter 5). Estimating a covariance matrix efficiently and understanding its graphical structure are difficult estimation problems because the number of unknown parameters in the covariance matrix increases quadratically with dimension and by the requirement that the estimate of the covariance matrix is positive definite.

There is a large literature of methods that use shrinkage or Bayesian models to improve on the maximum likelihood estimator of the covariance matrix. See, for example, Dempster (1969), Dempster (1972), Efron and Morris (1976), Yang and Berger (1994), Chiu et al. (1996), Giudici and Green (1999), Barnard et al. (2000), Wong et al. (2003) and Liechty et al. (2004). The simulation studies in Yang and Berger (1994) and Wong et al. (2003) show that considerable gains in efficiency are possible.

Dempster (1972) advocates a covariance selection approach to estimate a covariance matrix more efficiently, by which he means setting to zero some of the off-diagonal elements of the concentration matrix. His idea is that a more parsimonious model will give greater efficiency. However, the selection of which elements to set to zero is difficult even for moderate dimensions because a  $p \times p$  concentration matrix has  $p(p-1)/2$  distinct off-diagonal entries and there are  $2^{p(p-1)/2}$  possible graphs associated with it. Drton and Perlman (2004) give a model selection approach based on simultaneous confidence intervals to determine which partial correlations are zero. The simultaneous confidence intervals are based on large sample theory and become large when  $p$  is moderate to large. Drton and Perlman (2004) do not attempt to estimate the covariance matrix based on their selected graph.

A number of articles take a Bayesian approach to covariance selection. For the case of decomposable graphs, Dawid and Lauritzen (1993) introduces a conjugate prior for the covariance matrix called the hyper inverse Wishart distribution. Giudici (1996) uses a prior for the covariance matrix that is a mixture of fixed parameter hyper inverse Wishart priors over decomposable graphs and calculates the marginal likelihood for each decomposable graph. The marginal likelihood is used to calculate the posterior probability of each graph. This gives an exact solution for small examples, but for  $p$  greater than approximately 8 the number of graphs is prohibitively large.

Roverato (2000) shows that the hyper inverse Wishart prior for the covariance matrix is equivalent to a constrained Wishart prior for the concentration matrix. It is straightforward to define a constrained Wishart prior for general graphs, however, such distributions have normalizing constants that are not available analytically unless the graph is decomposable. Roverato (2002), Dellaportas et al. (2004) and Atay-Kayis and Massam (2005) propose efficient simulation and importance sampling methods for estimating the normalizing constants for the nondecomposable graphs. The normalizing constants are used to examine a small

number of graphs and select those that have the highest marginal likelihood or posterior probability, rather than to estimate the covariance matrix by averaging over graphs. However, such an approach seems unsuitable as the basis of a Markov chain Monte Carlo sampling scheme when  $p$  is moderate to large because there are  $2^{p(p-1)/2}$  possible graphs with only a small fraction of them being decomposable.

Giudici and Green (1999) give a MCMC approach that can deal with large values of  $p$ . Their method applies to a hierarchical model with a hyper inverse Wishart prior for the covariance matrix conditional on a decomposable graph. They use reversible jump Metropolis-Hastings methods to generate the covariance matrix and other parameters. Their method has a local computation property that only requires Cholesky decompositions of the submatrix of the covariance matrix corresponding to a clique of the graph. Brooks et al. (2003) modify the reversible jump MCMC proposal of Giudici and Green (1999) and give empirical results to show this improves the convergence rate.

Wong et al. (2003) also use MCMC methods to select which off-diagonal element to set to zero. They use reversible jump Metropolis-Hastings methods to generate the inverse covariance matrix and other parameters. The main difference between Giudici and Green (1999) and Wong et al. (2003) is that Wong et al. (2003) do not constrain the possible graphs to be decomposable. Wong et al. (2003) use a prior with normalizing constants based on graph size to avoid having to calculate normalizing constants for each nondecomposable graph. They also need to run a separate MCMC to estimate the normalizing constants for each graph size.

For longitudinal data, Smith and Kohn (2002) factor the concentration matrix using a Cholesky decomposition and carry out variable selection on the strict lower triangle of the Cholesky to obtain parsimony. Their approach is attractive when there is some natural ordering of the observation vector, but there are two potential drawbacks to the Cholesky approach when such a natural ordering does not exist. First, different orderings of the variables can yield different estimates of the covariance matrix. Second, under some orderings the Cholesky factor may be quite full even if the concentration matrix is sparse.

In this paper we consider Bayesian estimation of decomposable covariance selection models, also known as decomposable graphical Gaussian models. Our article makes the following contributions. First, we propose a prior for the covariance matrix such that the probability of each graph size is specified by the user, whereas most previous approaches, e.g. Giudici and Green (1999), assume that all graphs are equally probable. We show by simulation that the prior that assigns equal probability over graph sizes outperforms the prior that assigns equal probability over all graphs, both in identifying the correct decomposable model and in estimating the covariance matrix more efficiently. This advantage is greatest when the number of observations is small relative to the dimension of the covariance matrix. We also show by simulation that there is minimal change in statistical efficiency in using our mixture prior compared to the estimator of Wong et al. (2003), even when the graph of the covariance matrix is nondecomposable.

Our prior requires knowing the number of decomposable graphs for each graph size. The second contribution of the article is to give a MCMC method for estimating these counts, and to show that the counts obtained by the simulation method agree with analytic results when such results are known.

Our third contribution is to use the marginal likelihood results in Giudici (1996) to derive

a reduced conditional MCMC sampler for decomposable graphical models, where the covariance matrix is integrated out of all conditional distributions and is not generated in the MCMC. Our approach does not require reversible jump Metropolis-Hasting methods and has the local computation properties of the Giudici and Green (1999) approach, so the computational complexity for one iteration of our approach is similar to that of Giudici and Green (1999). We expect that our sampler has a faster convergence rate to that of Giudici and Green (1999) and Brooks et al. (2003), but we have not compared them empirically. We show that our sampler has a faster convergence rate than the Wong et al. (2003) approach. A comparison between the MCMC methods described in this paper and a stochastic search approach for finding the graph with maximum posterior probability is the subject of a forthcoming paper (Jones et al. (2005)).

The results in our article suggest that at present there is no ‘best’ method for estimating Gaussian covariance selection models. While the method of Wong et al. (2003) works in principle for all graphs, the convergence of their MCMC simulation can be slow if the true graph has full subgraphs of size 5 or larger because Wong et al. (2003) generate the elements of the concentration matrix one at a time. On the other hand the sampling scheme for decomposable graphs presented in our article is extremely efficient because the concentration matrix is integrated out and is an attractive alternative to the Wong et al. (2003) model for high dimensional graphs that are likely to have substantial full subgraphs. There are two other advantages of the decomposable prior considered in our article. The first is that there is a separate normalizing constant for each decomposable graph, whereas Wong et al. (2003) have a normalizing constant for each graph size. The second is that the Wong et al. (2003) model does not at present allow for hyperparameters in the prior. For example, using an equicorrelated prior as in Giudici and Green (1999) is not at present feasible with the approach of Wong et al. (2003).

The paper is organized as follows. Section 2 briefly introduces graphical Gaussian models. Section 3 describes our Bayesian covariance selection model and Section 4 describes our MCMC approach to estimating this model. Section 5 compares the prior that assigns equal probability to each graph size to the prior that assigns equal probability to each decomposable graph. Section 6 shows how to estimate the number of decomposable graphs for each size by simulation. Section 7 compares the prior that assigns equal probability to each graph size to the prior of Wong et al. (2003). There are two appendices. The first gives the proofs of the results in the paper. The second gives a computationally efficient expression for evaluating the ratio of normalizing constants from Section 4.1.

## 2. Background on Gaussian graphical models

Before explaining our Bayesian covariance selection model we provide some background on Gaussian graphical models. Further details on such models are available in Dawid and Lauritzen (1993) and Chapters 2, 3 and 5 of Lauritzen (1996).

Let  $g = (V, E)$  be an undirected graph with vertices  $V = \{1, \dots, p\}$  and set of edges  $E \subseteq V \times V$ . For a square matrix  $A$  we write  $A > 0$  to denote that  $A$  is positive definite. Let  $M^+(g)$  be the set of  $p \times p$  matrices  $\Omega$  satisfying  $\Omega > 0$  and  $\Omega_{ij} = 0$  for all pairs  $(i, j) \notin E$ .

For a given  $p \times p$  covariance matrix  $\Sigma$ , we define the *graph* of  $\Sigma$ ,  $g = g(\Sigma) = (V, E)$ , as follows. Let  $\Omega = \Sigma^{-1}$ . Let  $V = \{1, \dots, p\}$  and define  $E = \{(i, j), i \neq j \text{ such that } \Omega_{ij} \neq 0\}$ .

Thus the graph  $g = g(\Sigma)$  gives the configuration of nonzero off-diagonal elements in  $\Omega$ .

We say that an  $m \times m$  matrix  $A > 0$  has an inverse Wishart (IW) density with  $\delta > 0$  degrees of freedom and scale matrix  $\Phi$ , denoted as  $A \sim IW(m, \delta, \Phi)$ , if the density of  $A$  is

$$p(A|\delta, \Phi) = \frac{|\Phi|^{-\frac{\delta}{2}}}{\Gamma_m(\frac{\delta}{2})} |A|^{-\frac{(\delta+m+1)}{2}} \text{etr} \left( -\frac{1}{2} \Phi A^{-1} \right), \quad (1)$$

where  $\text{etr}(A) = \exp(\text{trace}(A))$  and for  $\alpha > \frac{(m-1)}{2}$ ,

$$\Gamma_m(\alpha) = \pi^{m(m-1)/4} \prod_{i=1}^m \Gamma(\alpha - (\frac{i-1}{2}))$$

is the multivariate gamma function (Muirhead, 1982, p. 113).

Lauritzen (1996, Definition 2.3, p.8) defines a *decomposable* graph and we refer to a covariance matrix  $\Sigma$  as decomposable if its graph  $g = g(\Sigma)$  is decomposable.

Suppose that  $g$  is a decomposable graph and let  $C_1, \dots, C_k$  be a perfect sequence of the cliques of  $g$ . Let  $H_j = C_1 \cup \dots \cup C_j$  be the history of the sequence and let  $S_j = H_{j-1} \cap C_j$  be the separators for  $j = 2, \dots, k$ . For any matrix  $M$  and subset of vertices  $B$ , use  $M_{BB}$  to denote the symmetric submatrix of  $M$  which is formed by taking every corresponding entry  $M_{ij}$  for which the vertices  $\{V_i, V_j\} \in B$ . Using the parameterization of Dawid (1981), we say  $\Sigma$  has a *hyper inverse Wishart* (HIW) distribution, with *hyperparameters*  $(\delta, \Phi)$  denoted by  $\Sigma \sim HIW(g, \delta, \Phi)$ , if for  $\Sigma^{-1} \in M^+(g)$

$$p(\Sigma|\delta, \Phi, g) = \frac{\prod_{i=1}^k p(\Sigma_{C_i C_i}|\delta, \Phi_{C_i C_i})}{\prod_{i=2}^k p(\Sigma_{S_i S_i}|\delta, \Phi_{S_i S_i})}, \quad (2)$$

where  $\delta > 0$ ,  $\Phi > 0$ , and the density is with respect Lebesgue measure on the elements of  $\Sigma$  corresponding to edges of  $g$ .

In (2), the terms  $p(\Sigma_{C_i C_i}|\delta, \Phi_{C_i C_i})$  denote the IW densities  $\Sigma_{C_i C_i} \sim IW(|C_i|, \delta + |C_i| - 1, \Phi_{C_i C_i})$  given by

$$p(\Sigma_{C_i C_i}|\delta, \Phi_{C_i C_i}) = \frac{|\Phi_{C_i C_i}|^{-\left(\frac{\delta+|C_i|-1}{2}\right)}}{\Gamma_{|C_i|}\left(\frac{\delta+|C_i|-1}{2}\right)} |\Sigma_{C_i C_i}|^{-\left(\frac{\delta+2|C_i|}{2}\right)} \text{etr} \left[ -\frac{1}{2} (\Sigma_{C_i C_i})^{-1} \Phi_{C_i C_i} \right], \quad (3)$$

where  $|C_i|$  denotes the cardinality of the clique  $C_i$ , and the terms  $p(\Sigma_{S_i S_i}|\delta, \Phi_{S_i S_i})$  are defined similarly. Note that the expression in (2) is invariant to the choice of perfect sequence.

From (1) – (3), the normalizing constant for the HIW distribution is

$$h(g, \delta, \Phi) = \frac{\prod_{i=1}^k \left[ \left| \frac{\Phi_{C_i C_i}}{2} \right|^{\left(\frac{\delta+|C_i|-1}{2}\right)} \Gamma_{|C_i|} \left( \frac{\delta+|C_i|-1}{2} \right)^{-1} \right]}{\prod_{i=2}^k \left[ \left| \frac{\Phi_{S_i S_i}}{2} \right|^{\left(\frac{\delta+|S_i|-1}{2}\right)} \Gamma_{|S_i|} \left( \frac{\delta+|S_i|-1}{2} \right)^{-1} \right]}. \quad (4)$$

### 3. Bayesian Covariance Selection models

#### 3.1. Likelihood and hierarchical structure

Suppose we have independent observations

$$y_t \sim N(\mu, \Sigma), \quad t = 1, \dots, n, \quad (5)$$

where  $y_t$  is  $p \times 1$ . Let  $y = (y_1, \dots, y_n)$  be the data. We use a hierarchical prior for  $\mu$  and  $\Sigma$  of the form

$$p(\mu, \Sigma, \Phi, \delta, g) = p(\mu|\Sigma, \Phi, \delta, g)p(\Sigma|\Phi, \delta, g)p(\Phi|\delta, g)p(\delta|g)p(g),$$

where each of the terms on the right is discussed below. In our article we assume that  $p(\mu|\Sigma, \Phi, \delta, g) \propto \text{constant}$ , as our focus is on priors for  $\Sigma$ . The prior for  $\Sigma$  depends on its graph  $g$ , the  $p \times p$  matrix  $\Phi$  and the scalar  $\delta$ , and is discussed in Section 3.2. Section 2 defines the graph of  $\Sigma$  as the configuration of nonzero off diagonal elements in  $\Sigma^{-1}$ . The prior for  $\Phi$  is discussed in Section 3.3 and the prior for the graph  $g$  is discussed in Section 3.4.

In the article we restrict the graph of  $\Sigma$  to be decomposable, so that the prior for  $\Sigma$  is a mixture over all decomposable graphs. We explain in Section 2 that this is equivalent to the prior for  $\Omega = \Sigma^{-1}$  being a mixture over all Wishart distributions constrained to decomposable graphs.

#### 3.2. Prior for $\Sigma$

We use the HIW prior (2) for  $\Sigma|\Phi, \delta, g$ , which allows  $\Sigma$  to be integrated out in the sampling scheme described in Section 4. Thus, our prior

$$p(d\Sigma|\Phi, \delta) = \sum_g p(d\Sigma|g, \Phi, \delta)p(g)$$

is a mixture of HIW distributions over all decomposable graphs  $g$ . As discussed in the introduction, Roverato (2000) shows that the inverse of a HIW random matrix has a Wishart distribution, subject to the constraints imposed by the corresponding graph. Thus

$$p(d\Sigma|\Phi, \delta) = \sum_g p(d\Sigma|g, \Phi, \delta)p(g)$$

is a mixture of of constrained Wishart distributions over all decomposable graphs.

In our article we set the degrees of freedom parameter  $\delta$  to 5 as such a value of  $\delta$  gives a suitably noninformative prior for  $\Sigma$ .

#### 3.3. Prior specification for $\Phi$ and its parameters

We consider the following three specifications for the hyperparameter  $\Phi$ , and refer to them as the hyperprior forms of  $\Phi$ :

- (a)  $\Phi = \tau I$ ,  $\tau > 0$  where  $I$  is the  $p \times p$  identity matrix.

- (b)  $\Phi = \tau(\rho J + (1 - \rho)I)$ ,  $\tau > 0$  where  $J$  is the  $p \times p$  matrix of ones and  $\rho$  is a correlation coefficient that needs to be in the open interval  $(-1/(p - 1), 1)$  for  $\Phi$  to be positive definite. This specification is used by Giudici and Green (1999) and is called the *equicorrelated version* of  $\Phi$  because  $\Phi_{ii} = \tau$  and  $\Phi_{ij} = \tau\rho$  for  $i \neq j$ .
- (c)  $\Phi = \tau S_y / (n - 1)$ , where  $\tau > 0$ ,

$$S_y = \sum_{t=1}^n (y_t - \bar{y})(y_t - \bar{y})', \quad (6)$$

and  $\bar{y}$  is the mean of the  $y_t$ .

We motivate the choice of  $\Phi$  in two ways. First, by integrating  $\mu$  out of  $p(y|\mu, \Sigma)$ , with  $p(\mu)$  constant, we obtain

$$p(y|\Sigma) \propto |\Sigma|^{-(n-1)/2} \text{etr} \left( -\frac{1}{2} S_y \Sigma^{-1} \right). \quad (7)$$

Suppose  $g$  is a decomposable graph. If we take  $p(\Sigma|g) \propto p(y|\Sigma)^{1/(n-1)}$ , then from (7) and equation (3) of Giudici (1996), we can write  $p(\Sigma|g)$  in the form (3) with  $\Phi = S_y / (n - 1)$ .

A second motivation for this choice of  $\Phi$  is to note that if  $\Sigma \sim HIW(p, \delta, \Phi)$ , then  $E(\Sigma_{CC}) = \Phi_{CC} / (\delta - 2)$  for any clique  $C = C_i$  or separator  $C = S_i$  in (2). Since  $(S_y)_{CC} / (n - 1)$  is an unbiased estimator of  $\Sigma_{CC}$ , this suggests taking  $\Phi \propto S_y / (n - 1)$ .

We assume in all cases that  $\tau$  is uniform on the interval  $[0, \Gamma]$  where  $\Gamma$  is large, e.g.  $\Gamma = 10^{10}$ , and in the equicorrelated case that  $\rho$  is uniform on the open interval  $(-1/(p - 1), 1)$ .

### 3.4. Prior for $g$

We first define notation for the *edge indicators* of a graph  $g$ . Let

$$e_{ij} = \begin{cases} 1 & \text{if } (i, j) \in E \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

and let  $e_{-ij} = \{e_{kl} : (k, l) \neq (i, j)\}$ . Note that any graph  $g = (V, E)$  can be unambiguously written as  $g = (e_{ij}, e_{-ij})$ .

For a given graph  $g = g(\Sigma)$ , let the number of edges, or the *size* of  $g$ , be given by

$$\text{size}(g) = \sum_{i < j} e_{ij} \quad (9)$$

i.e.  $\text{size}(g)$  is the number of nonzero elements in the strict upper triangle of  $\Omega$ , and  $\text{size}(g) \leq r = p(p - 1)/2$ .

Because of the theoretical and practical difficulty in calculating for a given  $p$  the exact number of decomposable graphs, or the number of graphs of a given size, most of the literature for both decomposable and general models takes the prior for  $g$  as uniform over all the relevant graphs; see, for example, Giudici and Green (1999), Dellaportas and Forster (1999), Geiger and Heckerman (2002), Giudici and Castelo (2003), Roverato (2002), and Atay-Kayis and Massam (2005).



Such a prior favours any class of graphs with many members over a class with few members, and favours middle sized graphs over both very large and very small sized graphs.

Let  $A_{p,k}$  denote the number of graphs of size  $k$ . We specify the prior for a graph  $g$  hierarchically as follows.

$$p(g|size(g) = k) = \frac{1}{A_{p,k}},$$

so that all graphs of a given size are equally likely. We now specify the prior for the size of a graph. One choice is

$$p(size = k) \propto A_{p,k},$$

which means that

$$p(g) = p(g, size(g)) = p(g|size(g))p(size(g)) \propto constant,$$

giving the uniform prior for  $g$ . A more flexible prior is of the form

$$p(size = k|\psi) = \binom{r}{k} \psi^k (1 - \psi)^{r-k},$$

where we interpret  $\psi$  as the probability that any two vertices have a common edge. We could then put a prior on  $\psi$ . Suppose we take the prior for  $\psi$  as a beta with parameters  $a$  and  $b$ , i.e.

$$p(\psi) = \frac{\psi^{a-1}(1-\psi)^{b-1}}{B(a,b)}.$$

Then,

$$p(size = k) = \binom{r}{k} \frac{B(a+k, r-k+b)}{B(a,b)}$$

and

$$p(g) = p(g|size(g))p(size(g)) \tag{10}$$

$$= \binom{r}{size(g)} \frac{B(a+size(g), r-size(g)+b)}{A_{p,size(g)}B(a,b)}, \tag{11}$$

where  $B(a,b)$  is the beta function. We could now also put a prior on  $a, b$ . In our article we take  $\psi$  uniform so that  $a = b = 1$ , which means that

$$p(size = k) = \frac{1}{(r+1)} \quad \text{and} \quad p(g) = \frac{1}{(r+1)A_{p,k}}.$$

That is, the size of each graph has equal probability, and the probability of a graph of size  $k$  conditional on  $size = k$  is uniform. However our framework is more flexible than this.

We call this the *size based prior* for  $g$  and compare results against those using a uniform prior.

The size based prior makes it easier to discover sparse and full graphs when  $n/p$  is small. The counts  $A_{p,k}$  are not available in the literature. Section 6 gives results to calculate a subset of them analytically, and shows how to evaluate the rest by simulation.

#### 4. Posterior inference and Markov chain Monte Carlo sampling

We use Markov Chain Monte Carlo (MCMC) simulation to obtain all posterior distributions. The simulation involves the generation of the graphs  $g$  and the parameters in  $\Phi$  but not  $\Sigma$  and  $\mu$  which are integrated out. Thus, our sampling scheme is said to generate from *reduced conditionals* and is therefore expected to be more efficient than the sampling schemes in Giudici and Green (1999) and Wong et al. (2003) that generate  $\Sigma$  as part of their sampling scheme.

We note that iterates of  $\mu$  and  $\Sigma$  can also be generated in conjunction with the simulation, but such iterates of  $\mu$  and  $\Sigma$  do not have any influence on the convergence properties or dependence structure of the reduced conditional simulation.

The following theorems are useful in evaluating the conditional distributions required in the simulations. The first theorem gives a conjugate prior property of the HIW distribution.

Let  $S_y$  be defined by (6) and define

$$\Phi^* = \Phi + S_y \text{ and } \delta^* = \delta + n - 1. \quad (12)$$

**THEOREM 1.** (Dawid and Lauritzen, 1993) For the Bayesian model specified by (2) and (7)

$$\Sigma|y, \delta, \Phi, g \sim HIW(g, \delta^*, \Phi^*).$$

**PROOF.** See Dawid and Lauritzen (1993) or Appendix A.

The next theorem gives an expression for the marginal likelihood.

**THEOREM 2.** (Giudici, 1996) For the Bayesian model specified by (2) and (7),

$$p(y|\delta, \Phi, g) = (2\pi)^{-((n-1)p/2)} \frac{h(g, \delta, \Phi)}{h(g, \delta^*, \Phi^*)} \quad (13)$$

**PROOF.** See Giudici (1996) or Appendix A.

##### 4.1. Sampling the graphs $g$

We sample the graphs  $g$  by generating the edge indicators one at a time, conditional on  $\delta, \Phi$  and  $e_{-ij} = \{e_{kl}, (k, l) \neq (i, j), k < l\}$  using the following MH sampling scheme.

Using the notation of Section 3, let  $g^c = (V, E^c)$  be the current graph of  $\Sigma$ , which is decomposable by construction with edge indicators  $\{e_{kl}^c : 1 \leq k < l \leq p\}$ .

We choose a pair  $(i, j)$  at random and suppose that  $g = (e_{ij}, e_{-ij}^c)$  is decomposable for both  $e_{ij} = 0$  and  $e_{ij} = 1$ . We use the legal edge addition and deletion characterizations of Giudici and Green (1999) and Frydenberg and Lauritzen (1989) respectively to ensure this. Otherwise we choose a new pair  $(i, j)$ .

Set the proposal graph as  $g^p$  (conditional on  $g^c$ ) as  $g = (e_{ij}^p, e_{-ij}^c)$  where  $e_{ij}^p = 1 - e_{ij}^c$ . This means that the proposal density for  $e_{ij}$  is  $q_g(a|b, e_{-ij}^c)$  where  $a$  and  $b$  are each either 0 or 1, and  $q_g(a = 1 - b|b, e_{-ij}^c) = 1$ .

The MH acceptance probability for the proposal is

$$\min \left\{ 1, \frac{p(y|g^p, \Phi, \delta) p(g^p)}{p(y|g^c, \Phi, \delta) p(g^c)} \right\} \quad (14)$$

because  $q_g(e_{ij}^c | e_{ij}^p, e_{-ij}^c) / q_g(e_{ij}^p | e_{ij}^c, e_{-ij}^c) = 1$ . The ratio  $p(g^p)/p(g^c)$  is known and the ratio of marginal likelihoods

$$\frac{p(y|g^p, \Phi, \delta)}{p(y|g^c, \Phi, \delta)} = \frac{h(g^p, \delta, \Phi) h(g^c, \delta^*, \Phi^*)}{h(g^c, \delta, \Phi) h(g^p, \delta^*, \Phi^*)}. \quad (15)$$

A simple expression for (15) is derived in Appendix B.

#### 4.2. Generating the parameters in $\Phi$

In all cases of Section 3.3 we generate  $\tau$  using a random walk MH method

$$\log(\tau^p) = \log(\tau^c) + \xi_\tau, \quad \xi_\tau \sim N(0, \sigma_\tau^2),$$

which has acceptance probability

$$\min \left\{ 1, \frac{p(y|g, \tau^p, \rho) p(\tau^p)}{p(y|g, \tau^c, \rho) p(\tau^c)} \right\} \quad (16)$$

as the proposal densities cancel out. In the equicorrelated case, the parameter  $\rho$  is generated similarly to  $\tau$  by a random walk MH method

$$\rho^p = \rho^c + \xi_\rho, \quad \xi_\rho \sim N(0, \sigma_\rho^2).$$

The choice of the variances  $\sigma_\tau^2, \sigma_\rho^2$  is sensitive to  $p$ , and was fine tuned to attain acceptance probabilities of around 25% according to the acceptance rate of the proposals. For the case  $p = 17$  reported in this paper, such an acceptance probability resulted from using  $\sigma_\tau^2 = 1/10$  and  $\sigma_\rho^2 = 1/20$ .

#### 4.3. Generating $\Sigma, \Omega$ and $\mu$

Although  $\mu, \Sigma$  and  $\Omega$  are not generated in the MCMC simulation, it is often necessary to estimate functionals of  $\mu, \Sigma$  and  $\Omega$ . Such functionals can be estimated by sampling from the posterior distribution of  $\Sigma, \Omega$  and  $\mu$ . Conditional on  $(g, \delta, \Phi)$  it follows from Theorem 1 that  $p(\Sigma|y, g, \delta, \Phi)$  is HIW  $(\delta + n - 1, \Phi + S_y)$  so that  $\Sigma$  and  $\Omega$  can be generated using Theorems 3 and 4 of Roverato (2000). It is straightforward to show that  $p(\mu|y, \Sigma, g, \delta, \Phi)$  is  $N(\bar{y}, \Sigma/n)$ , and hence to generate  $\mu$ , giving iterates  $\{\mu^{[j]}, \Sigma^{[j]}, \Omega^{[j]}, j \geq 1\}$  from the posterior distribution.

#### 4.4. Efficient estimation of $E(\Omega|y)$

The posterior mean of  $\Omega$  is not only used as an estimator of  $\Omega$ , but also of  $\Sigma$  because  $E(\Omega|y)^{-1}$  is the Bayes estimator of  $\Sigma$  for the  $L_1$  loss function in Section 5. One method

of estimating  $E(\Omega|y)$  is to use the histogram estimator  $J^{-1} \sum_{j=1}^J \Omega^{[j]}$ . A statistically more efficient estimator is the mixture estimator  $J^{-1} \sum_{j=1}^J E(\Omega|y, g^{[j]}, \delta^{[j]}, \Phi^{[j]})$ .

We now show how to efficiently compute  $E(\Omega|y, g, \delta, \Phi)$  using the following notation from Lauritzen (1996). Suppose that  $A$  is a  $p \times p$  matrix and  $S \subset V$ . Let  $B = [A_{SS}]^V$  be the  $p \times p$  matrix defined by

$$B_{ij} = \begin{cases} A_{ij} & \text{if } \{i, j\} \subset S \\ 0 & \text{otherwise} \end{cases}$$

**THEOREM 3.** *Suppose that  $\Omega|y \sim W(g, \delta^*, \Phi^*)$ , where  $g$  is decomposable. Then, using the notation of this and Section 3,*

$$E(\Omega|y, \delta, \Phi, g) = \sum_{i=1}^k [(\delta^* + |C_i| - 1) (\Phi_{C_i C_i}^*)^{-1}]^V - \sum_{i=2}^k [(\delta^* + |S_i| - 1) (\Phi_{S_i S_i}^*)^{-1}]^V. \quad (17)$$

**PROOF.** See Appendix A.

## 5. Comparison of the size prior for a graph with the uniform prior

This section compares the prior based on the graph size with the uniform prior that is used in most previous articles. Performance is in terms of a loss function and a simulation was carried out to numerically assess performance. We found that overall the size based prior for  $g$  outperformed the uniform prior.

Our simulation considered the following five graph types for  $g$ . (a)  $\Omega = I$ , the identity matrix, representing the empty graph and a diagonal covariance matrix; (b)  $\Omega$  tridiagonal, representing a sparse and decomposable graph (this is a chain graph with  $p - 1$  edges); (c)  $\Omega$  an ‘extreme’ full matrix (the correlation coefficients  $\rho_{ij}$  of  $\Omega^{-1}$  satisfy  $|\rho_{ij}| > .30$ ), which is a complete graph; (d)  $\Omega$  corresponding to a 4-cycle on  $p$  vertices representing a sparse but nondecomposable graph; and (e)  $\Omega$  corresponding to a  $p$ -cycle on  $p$  vertices, again representing a sparse but nondecomposable graph. We note that the nondecomposable graphs in (d) and (e) require the addition of extra edges when we estimate them by a mixture of decomposable graphs. Furthermore, (e) is an extreme case of non-decomposability, as it requires a minimum of  $p - 3$  fill ins. Conversely, the unchorded 4-cycle on  $p$  nodes requires the fewest number of fill ins, so was chosen as an indicator of performance for the sparsest nondecomposable case.

The simulation considered the three forms of  $\Phi$  described in Section 3.3 and two sample sizes  $n = 40$  and  $n = 100$ . We report results for matrices of size  $p = 17$ , but similar results were obtained for matrices of other sizes.

Let  $\Sigma_T$  be the true value of  $\Sigma$  and let  $\widehat{\Sigma}$  be an estimator of  $\Sigma_T$ . We measure the performance of  $\widehat{\Sigma}$  using the  $L_1$  loss function

$$L_1(\widehat{\Sigma}, \Sigma_T) = \text{trace}(\widehat{\Sigma}\Sigma_T^{-1}) - \log \det(\widehat{\Sigma}\Sigma_T^{-1}) - p. \quad (18)$$

This loss function is frequently used to compare estimates of the covariance matrix, e.g. Yang and Berger (1994). It is straightforward to show that  $L_1 \geq 0$  for all  $\widehat{\Sigma}$  and  $\Sigma_T$ , and that it

is only equal to 0 if  $\widehat{\Sigma} = \Sigma_T$ . It is also straightforward to show that for  $y \sim N(0, \Sigma)$ ,

$$L_1(\widehat{\Sigma}, \Sigma_T) = - \int p(y|\widehat{\Sigma}) \log \left( \frac{p(y|\Sigma_T)}{p(y|\widehat{\Sigma})} \right) dy \quad (19)$$

i.e.  $L_1$  is equivalent to a Kullback-Liebler distance between  $p(y|\Sigma_T)$  and  $p(y|\widehat{\Sigma})$  with respect to the density  $p(y|\widehat{\Sigma})$ . The Bayes estimator for  $\Sigma$  for the  $L_1$  loss function is  $E(\Omega|y)^{-1}$ , which can be computed as in Section 4.4.

We use boxplots to compare replication by replication the size based prior with the uniform prior in terms of the percentage *increase* in the loss function  $L_1$  resulting from using the uniform prior compared to the size-based prior; i.e. the boxplots are based on calculating

$$100(L_1^{unif} - L_1^{size})/L_1^{size}$$

for each replication, where  $L_1^{unif}$  and  $L_1^{size}$  are the values of  $L_1(\widehat{\Sigma}, \Sigma_T)$  for the uniform and size based priors respectively.

The boxplots are based on 20 replications with each replication consisting of 2,000 burnin iterations and 20,000 sampling iterations. We ran the sampler for the case  $p = 17$  on  $n = 40$  and 100 observations from five simulated data sets corresponding to the five models (a)–(e) for  $\Omega$ .

Figure 1 presents the results for  $p = 17$ . The plots show that for  $\Phi = \tau I$  and  $\Phi$  equicorrelated, the size prior is at least as good, and often much better than, the uniform prior. For  $\Phi = \tau S_y/(n - 1)$ , the comparison between the size prior and the uniform prior is inconclusive for  $n = 40$ , but for  $n = 100$  the size prior is at least as good as, and often better than the uniform prior. We conclude that the size based prior outperforms the uniform prior.

We also compared the performance of the three forms of  $\Phi$  for the uniform and size priors and found that overall the equicorrelated form of  $\Phi$  using the size based prior for the graph performed best, and it is this combination that we use for the rest of the paper.

## 6. Evaluating the size based prior

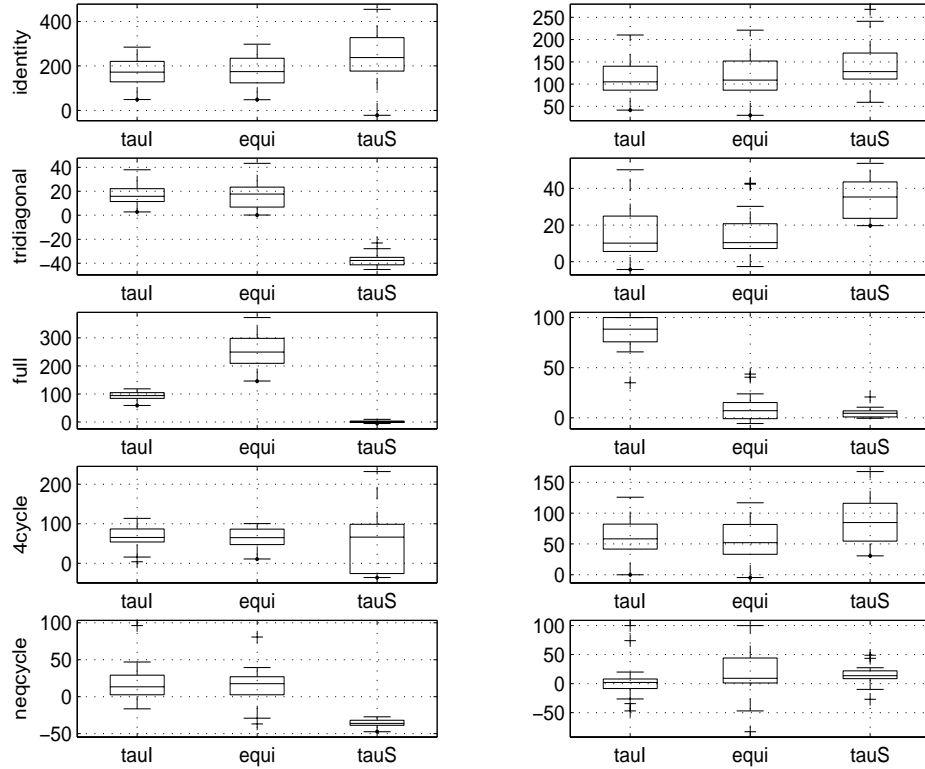
To use the size based prior for graphs on  $p$  vertices, we need the set of numbers  $\{A_{p,k} : k = 0, \dots, r\}$  where  $A_{p,k}$  is the number of decomposable graphs of size  $k$  on  $p$  vertices, and  $r = \binom{p}{2}$  is the maximum graph size. These numbers are not in the literature, nor is there a general method available for computing them. In this section we present some exact values of  $A_{p,k}$  as well as a simulation method that can estimate the  $A_{p,k}$  as precisely as necessary.

Let  $B_{p,k}$  be the number of connected decomposable graphs of size  $k$  on  $p$  vertices. Equations (3) and (4) of Castelo and Wormald (2001) give recurrences to calculate  $A_{p,k}$  from the  $B_{p,k}$  analytically, and the information to calculate all  $B_{p,k}$  analytically is implicit in Wormald (1985). For  $p \leq 8$ , Wormald (1985) gives the  $B_{p,k}$  from which we computed the  $A_{p,k}$  and these are reported in Table 1.

However, Wormald's (1985) analytic approach for obtaining the  $B_{p,k}$  is likely to be computationally intractable for  $p > 25$  (private correspondence with Wormald) and even for  $8 < p \leq 25$  obtaining the  $B_{p,k}$  would take weeks on realistically sized computers. Furthermore, analytically deriving the  $A_{p,k}$  from the  $B_{p,k}$  is computationally feasible only for small

**Table 1.** For each  $p$ ,  $2 \leq p \leq 8$  the table gives each  $A_{p,k}$ ,  $0 \leq k \leq r$  and  $A_p = \sum_{k=0}^r A_{p,k}$ . The table also gives for each  $p$  the percentage of graphs that are decomposable.

k	2	3	4	5	6	7	8
0	1	1	1	1	1	1	1
1	1	3	6	10	15	21	28
2		3	15	45	105	210	378
3		1	20	120	455	1330	3276
4			12	195	1320	5880	20265
5			6	180	2526	18522	92988
6			1	140	3085	40647	315574
7				90	3255	60795	770064
8				30	3000	79170	1357818
9				10	2235	92785	2078300
10				1	1206	94521	2892176
11					615	81417	3621576
12					260	58485	4016439
13					60	40110	3916724
14					15	24255	3432660
15					1	12222	2855748
16						4872	2185484
17						1890	1488984
18						595	902944
19						105	493220
20						21	258468
21						1	118504
22							46046
23							14868
24							4690
25							1176
26							168
27							28
28							1
$\sum_{k=0}^r A_{p,k}$	2	8	61	822	18,154	617,675	30,888,596
% decomposable	100%	100%	95%	80%	55%	29%	12%



**Fig. 1.** Percentage increase in loss of uniform prior relative to the size prior measured under  $L_1$  loss. The left panels correspond to  $n = 40$  and the right panels to  $n = 100$ . taul, equi and tauS correspond to  $\Phi = \tau I$ ,  $\Phi$  equicorrelated and  $\Phi = \tau S_y / (n - 1)$ .

$p$ . Because of these difficulties we propose a simulation methodology to estimate the  $A_{p,k}$  for all  $p$ .

### 6.1. Methodology

We begin with some exact results which can be used to calculate  $\{A_{p,k} : k \leq 5 \text{ and } r - 2 \leq k \leq r\}$  analytically for any  $p$ . Let  $F_{p,k}$  denote the number of nondecomposable graphs having  $p$  vertices and  $k$  edges.

- LEMMA 4. (a)  $A_{p,k} = \binom{r}{k} - F_{p,k}$ .  
 (b)  $F_{p,0} = F_{p,1} = F_{p,r} = 0$ ,  $p \geq 0$ .  
 (c)  $F_{p,2} = F_{p,r-1} = 0$ ,  $p \geq 2$ .  
 (d)  $F_{p,3} = 0$ ,  $p \geq 3$ .

PROOF. The proof is obvious.

- LEMMA 5. (a) For  $p \geq 4$ ,  $F_{p,4} = \binom{p}{4} \times 3$ .  
 (b) For  $p \geq 4$ ,  $F_{p,r-2} = F_{p,4}$ .  
 (c) For  $p \geq 5$ ,  $F_{p,5} = \binom{p}{5} \times 12 + \binom{p}{4} \times 3 \times (r - 6)$ .

PROOF. See Appendix A.

We now show how to estimate the  $\{A_{p,k} : 6 \leq k \leq r - 3\}$  for all  $p$ . Our approach is to run a separate simulation to estimate each  $A_{p,k}$  for  $6 \leq k \leq r - 3$ . The simulations are done in ascending order of  $k$ , i.e.  $k = 6, \dots, r - 3$ , and the simulation to estimate a particular  $A_{p,k}$  is restricted to graphs of size  $\leq k$  and uses the estimates  $\widehat{A}_{p,j}$  of  $A_{p,j}$  for  $j = 6, \dots, k - 1$  that have been calculated in previous simulations.

We now describe the details of the simulation to estimate a particular  $A_{p,k}$ . Let  $\phi_{p,k}$  be the initial estimate of  $A_{p,k}$  given by

$$\phi_{p,k} = \widetilde{\alpha}_{p,k} \frac{\widehat{A}_{p,k-1}^2}{\widehat{A}_{p,k-2}} \quad (20)$$

with  $\widetilde{\alpha}_{p,k}$  chosen in the range  $(0.5, 1)$ . To justify this choice of  $\phi_{p,k}$ , we note that we have found empirically that  $\log A_{p,k}$  is approximately a negative quadratic (see figures 2 and 3) so that  $\log A_{p,k} - 2 \log A_{p,k-1} + \log A_{p,k-2} \leq 0$ , and hence

$$\alpha_{p,k} = \frac{A_{p,k}/A_{p,k-1}}{A_{p,k-1}/A_{p,k-2}} \leq 1.$$

We have also found empirically that  $\alpha_{p,k}$  is likely to exceed 0.5.

As

$$A_{p,k} = \alpha_{p,k} \frac{A_{p,k-1}^2}{A_{p,k-2}}$$

the above discussion suggests the choice of  $\phi_{p,k}$  in (20). Further details on the choice of  $\widetilde{\alpha}_{p,k}$  can be obtained from the authors.

We use Lemmas 4 and 5, the estimates  $\widehat{A}_{p,j}$  of  $A_{p,j}$  for  $j = 6, \dots, k - 1$  that have been calculated in previous simulations, and the initial estimate  $\phi_{p,k}$  of  $A_{p,k}$  given above to define the following probability distribution  $p_e(g)$  on the graphs  $g$  of size  $\leq k$ . To simplify the notation we omit subscripts for  $p$  and  $k$  in  $p_e(g)$ .

$$p_e(g) \propto \begin{cases} \frac{1}{A_{p, \text{size}(g)}} & \text{if } 0 \leq \text{size}(g) \leq 5 \\ \frac{1}{A_{p, \text{size}(g)}} & \text{if } 6 \leq \text{size}(g) \leq k - 1 \\ \frac{1}{\phi_{p,k}} & \text{if } \text{size}(g) = k \end{cases} \quad (21)$$

which implies that

$$\begin{aligned} \frac{p_e(\text{size} = k)}{p_e(\text{size} \leq 5)} &= \frac{A_{p,k}/\phi_{p,k}}{\sum_{j=0}^5 A_{p,j}/A_{p,j}} \\ &= \frac{1}{6} A_{p,k}/\phi_{p,k} \end{aligned}$$



and hence

$$A_{p,k} = 6\phi_{p,k} \frac{p_e(\text{size} = k)}{p_e(\text{size} \leq 5)}.$$

By running the simulation described below based on  $p_e(g)$  we can estimate the ratio  $p_e(\text{size} = k)/p_e(\text{size} \leq 5)$  by their relative frequencies and hence obtain an estimate of

$$\widehat{A}_{p,k} = 6\phi_{p,k} \frac{\widehat{p}_e(\text{size} = k)}{\widehat{p}_e(\text{size} \leq 5)},$$

where  $\widehat{p}_e(\text{size} = k)$  and  $\widehat{p}_e(\text{size} \leq 5)$  are the empirical relative frequencies.

The simulation uses the following MCMC sampling scheme. As in Section 4.1, we generate the edge indicators one at a time conditional on the other edge indicators. Let  $g^c = (V, E^c)$  be the current graph with edge indicators given by  $\{e_{kl} : (k, l) \in E^c\}$ . We select an edge  $(i, j)$  at random. If  $g = (e_{ij}, e_{-ij}^c)$  corresponds to a decomposable graph of size  $\leq k$  for both  $e_{ij} = 0$  and  $e_{ij} = 1$  then we proceed, where we again use the legal edge addition and deletion characterizations of Giudici and Green (1999) and Frydenberg and Lauritzen (1989) respectively to test this. Otherwise we select a new edge. If we proceed, then we propose a new graph  $g^p = (1 - e_{ij}^c, e_{-ij}^c)$  and accept this graph with probability

$$\min\{1, p_e(g^p)/p_e(g^c)\}$$

which is evaluated using (21).

We note that at each stage we can also re-estimate  $A_{p,j}$ ,  $j = 6, \dots, k - 1$ .

## 6.2. Results


This section presents the estimates  $\widehat{A}_{p,k}$  for  $k = 0 \dots r$  and  $p = 8$  and 34. We also provide a general method to check on the quality of these estimates. We note that the prior

$$p_e(g) \propto \begin{cases} \frac{1}{A_{p, \text{size}(g)}} & \text{if } 0 \leq \text{size}(g) \leq 5 \text{ or } r - 2 \leq \text{size}(g) \leq r \\ \frac{1}{\widehat{A}_{p, \text{size}(g)}} & \text{if } 6 \leq \text{size}(g) \leq r - 3 \end{cases}$$


should have  $p_e(\text{size} = k)$  close to uniform and hence close to the target value of  $1/(r+1)$ , and that an approximate lower bound for the standard error of the estimates of  $p_e(\text{size} = k)$  is  $\sqrt{\pi(1-\pi)/J}$ , where  $\pi = 1/(r+1)$  and  $J$  is the number of iterates used to compute  $p_e(\text{size} = k)$ . Our simulations use a burnin period of 2,000 iterations and a sampling period of  $N = 10,000$  iterations. Figure 2 plots the estimates  $\widehat{A}_{p,k}$  for  $p = 8$  and the true values  $A_{8,k}$ ,  $k = 0 \dots r$  on both an absolute and logarithmic scale. Figure 2 also plots the estimates of  $p_e(\text{size} = k)$  together with the target value  $1/(r+1)$  and lower bounds for the  $\pm 3$  standard error lines.

Figure 3 has the same interpretation as Figure 2 but is for  $p = 34$ . The true values of  $A_{34,k}$  are not plotted as they are mostly unknown.

For  $p = 9, \dots, 12$  the totals  $A_p = \sum_j A_{p,j}$  are known, but not the  $A_{p,j}$ . As a further check on results we compared our estimated values of  $\widehat{A}_p$  to  $A_p$  and found that we were consistently within 1% of the truth.



**Fig. 2.** Panel (a): Plot of true  $A_{S,k}(\cdot)$  and estimates  $\hat{A}_{S,k}$  (open circles),  $k = 0, \dots, r$ . Panel (b): Log scale of plot (a). Panel (c): Plot of  $\hat{p}_e(\text{size} = k)$  together with their target value of  $1/(r + 1)$  (middle horizontal line) and  $\pm 3$  approximate standard errors (outer horizontal lines).



**Fig. 3.** Panel (a): Plot of estimates  $\hat{A}_{34,k}$   $k = 0, \dots, r$ . Panel (b): Log scale of plot (a). Panel (c): Plot of  $\hat{p}_e(\text{size} = k)$  together with their target value of  $1/(r + 1)$  (middle horizontal line) and  $\pm 3$  approximate standard errors (outer horizontal lines).

## 7. Comparison to the Wong et al. (2003) covariance selection prior

This section compares the performance of the prior in our article to the covariance selection prior of Wong et al. (2003), which does not assume that the graph of the covariance matrix is decomposable. Based on the results in Section 5, we use the equicorrelated form of  $\Phi$  and the size based prior for the decomposable graphs.

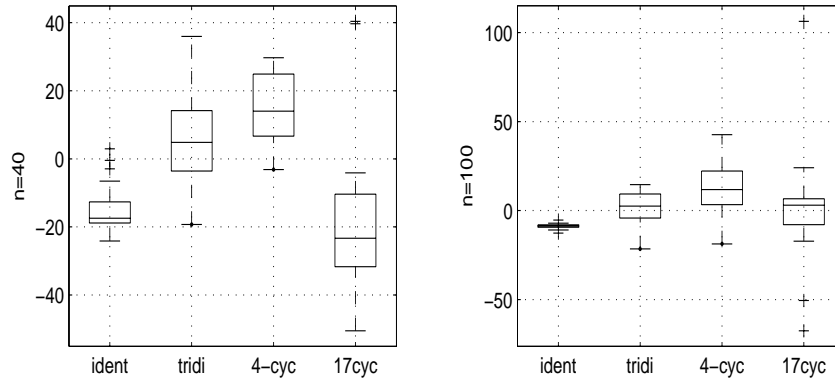
The design of the simulation study is similar to that in Section 5. We use  $L_1$  as the loss function,  $p = 17$ , two sample sizes  $n = 40$  and  $n = 100$ , and four graphs for  $\Omega$ : identity, tridiagonal, 4-cycle and 17-cycle.

We refer to the decomposable prior as *DCP* and the nondecomposable prior of Wong et al. (2003) as *NDP*. Figure 4 reports boxplots of the percentage increase in  $L_1$  of *DCP* over

$NDP$  for each iterate, i.e.

$$100(L_1^{DCP} - L_1^{NDP})/L_1^{NDP}.$$

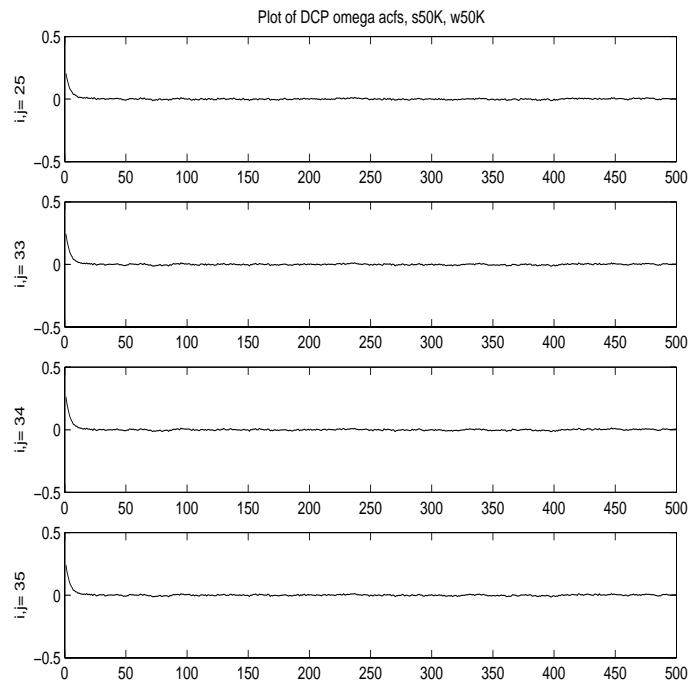
Figure 4 shows that both priors perform similarly for decomposable graphs and nondecomposable graphs, for both  $n = 40$  and  $n = 100$ . These results and others suggest that the prior based on decomposable graphs performs similarly to that of Wong et al. (2003) when the graphs are relatively sparse.




**Fig. 4.** Percentage increase in  $L_1$  for  $DCP$  over  $NDP$ . The left panel is for  $n = 40$  and the right is for  $n = 100$ .

Next we report autocorrelation plots for the iterates of the elements of  $\Omega$ , when  $p = 5$  and the graph is full for both  $DCP$  and  $NDP$  when  $n = 40$ . The simulation for  $DCP$  uses a burnin of 50,000 iterations and a sampling of 50,000 iterations, and 500,000 burnin and 1 million sampling iterations for  $NDP$ .

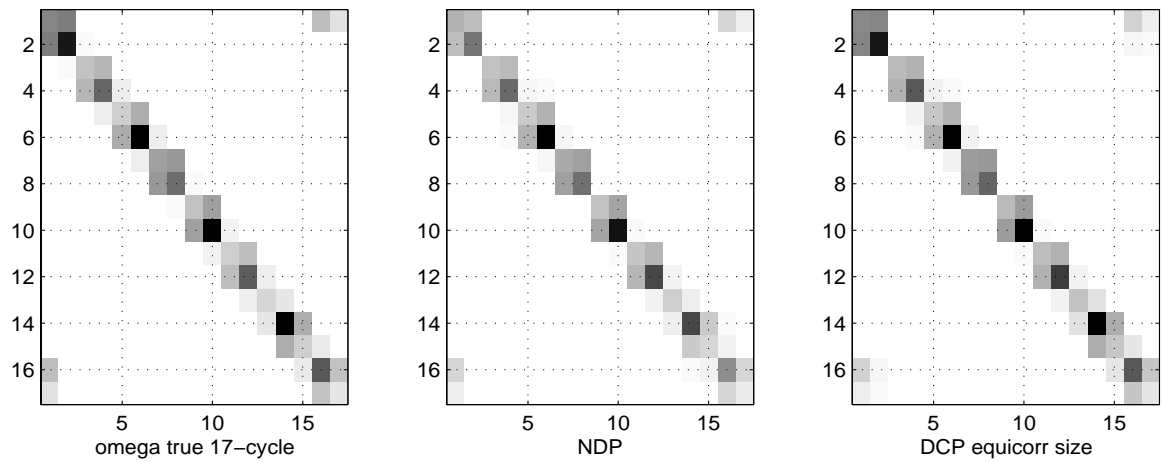
Figures 5 and Figure 6 are the autocorrelation plots for the  $DCP$  and  $NDP$  models for a representative selection of  $\Omega_{ij}$ . The figures show that the autocorrelations of the iterates of the  $\Omega_{ij}$  decay rapidly to zero for the  $DCP$  model, but are far more dependent in the  $NDP$  model. This difference in dependence is due to the greater efficiency of the sampling scheme in the decomposable case. Grey scale plots of the true inverse covariance  $\Omega$  and posterior mean estimates of  $\Omega$  for the  $NDP$  estimator and the  $DCP$  estimator for the 17-cycle case indicated that  $NDP$  and  $DCP$  performed similarly in the simulations. For brevity only the nondecomposable 17-cycle is presented as it represents a case of high non-decomposability. Figure 7 shows that even in this case, the grey scales are very similar.



**Fig. 5.** Autocorrelations of the iterates of the  $\Omega_{i,j}$  in the *DCP* case for a representative selection of  $\Omega_{i,j}$ .



**Fig. 6.** Autocorrelations of the iterates of the  $\Omega_{ij}$  in the *NDP* case for a representative selection of  $\Omega_{ij}$ .



**Fig. 7.** True inverse covariance  $\Omega$  and posterior mean estimates of  $\Omega$  for the *NDP* estimator and the *DCP* estimator for the 17-cycle case.

## A. Proofs of results

### Proof of Theorem 1

Roverato (2000) shows that if  $\Sigma \sim HIW(g, \delta, \Phi)$  and  $\Omega = \Sigma^{-1}$  then

$$p(\Omega|g, \delta, \Phi) \propto |\Omega|^{(\delta-2)/2} \text{etr} \left( -\frac{1}{2} \Omega \Phi \right) \quad (22)$$

The result then follows from (7) since

$$\begin{aligned} p(\Omega|y, g, \delta, \Phi) &\propto p(y|\Omega)p(\Omega|g, \delta, \Phi) \\ &\propto |\Omega|^{(n-1)/2} \text{etr} \left( -\frac{1}{2} \Omega S_y \right) |\Omega|^{(\delta-2)/2} \text{etr} \left( -\frac{1}{2} \Omega \Phi \right) \\ &= |\Omega|^{(n+\delta-3)/2} \text{etr} \left( -\frac{1}{2} \Omega (S_y + \Phi) \right). \end{aligned}$$

Note that the conjugate prior result for  $\Omega$  does not require the graph  $g$  to be decomposable.

### Proof of Theorem 2

First

$$p(Y|\delta, \Phi, g) = \frac{p(Y|\Sigma, \delta, \Phi, g)p(\Sigma|\delta, \Phi, g)}{p(\Sigma|Y, \delta, \Phi, g)}.$$

The result then follows from (2), (3), (7) and Theorem 1.

### Proof of Theorem 3

From Equation (5.23), Lemma 5.5 of Lauritzen (1996)

$$\Omega = \sum_{i=1}^k \left[ (\Sigma_{C_i C_i})^{-1} \right]^V - \sum_{i=2}^k \left[ (\Sigma_{S_i S_i})^{-1} \right]^V$$

and hence

$$E(\Omega|Y, \delta, \Phi, g) = \sum_{i=1}^k \left[ E \left( (\Sigma_{C_i C_i})^{-1} | Y, \delta, \Phi, g \right) \right]^V - \sum_{i=2}^k \left[ E \left( (\Sigma_{S_i S_i})^{-1} | Y, \delta, \Phi, g \right) \right]^V.$$

Now  $\Sigma|Y, \delta, \Phi, g \sim HIW(\delta, \Phi^*, g^*)$ , so from Dawid and Lauritzen (1993), if  $A$  is a complete set in  $g$  then  $(\Sigma_{AA})^{-1} | Y, \delta, \Phi, g \sim \text{Wishart}(\delta^* + |A| - 1, \Phi_{AA}^*)$ . The result then follows from the properties of the Wishart distribution.

### Proof of Lemma 5

- (a) For a nondecomposable graph to have 4 edges it must contain exactly one chordless 4-cycle and no other edges. There are  $\binom{p}{4}$  possible choices for the 4 vertices, and for each choice of 4 vertices there are 3 different chordless 4-cycles.
- (b) For a graph to be nondecomposable with  $\binom{p}{2} - 2$  edges it must contain exactly one 4 cycle and all other edges must be present. Then apply the proof of the above.
- (c) We can partition the nondecomposable graphs with 5 edges into 2 sets: (a) those with a chordless 5-cycle and no other edges, and (b) those with a chordless 4-cycle and an extra edge. For case (a) there are  $\binom{p}{5}$  choices for the 5 vertices and for each choice there are  $(5-1)!/2 = 12$  different chordless 5-cycles. For case (b) there are  $\binom{p}{4} \times 3$  choices for the chordless 4-cycle, and for each choice of chordless 4-cycle there are  $(\binom{p}{2} - 6)$  choices for the extra vertex pair constituting the edge.

## B. HIW results for Bayesian analysis using MCMC

The following results derive an expression for (15) that can be evaluated efficiently. The first theorem gives some necessary graph theory.

Let  $g = (V, E)$  be a decomposable graph with edge indicators  $\{e_{ij}, i < j \leq p\}$ . Assume the edge indicator  $e_{ij} = 1$  for  $g$ , and that the graph  $g' = (V, E')$  is decomposable and has edge set  $E'$  as defined by indicators  $\{e'_{ij} = 0, e_{-ij}\}$ .

**THEOREM 6.** *Suppose that  $g$  and  $g'$  are the decomposable graphs defined above. Suppose that  $C_1, \dots, C_k$  are the cliques of  $g$  ordered to form a perfect sequence and  $S_2, \dots, S_k$  are the corresponding separators. Then*

- (a) *The edge  $(i, j)$  is contained in a single clique of  $g$ .*
- (b) *If  $(i, j) \in C_q$  then either  $i \notin S_q$  or  $j \notin S_q$ .*
- (c) *If  $j \notin S_q$  and  $C_{q_1} = C_q \setminus \{j\}$  and  $C_{q_2} = C_q \setminus \{i\}$  then  $C_1, \dots, C_{q-1}, C_{q_1}, C_{q_2}, C_{q+1}, \dots, C_k$  is a perfect sequence of complete sets in  $g'$  and has separators  $S_2, \dots, S_{q-1}, S_{q_1} = S_q, S_{q_2} = C_q \setminus \{i, j\}, S_{q+1}, \dots, S_k$ .*
- (d) *The sequence  $C_1, \dots, C_{q-1}, C_{q_1}, C_{q_2}, C_{q+1}, \dots, C_k$  contains all the cliques of  $g'$ .*

**PROOF.** Part (a) is Theorem 1 of Frydenberg and Lauritzen (1989).

Parts (b) and (c) follow from part (a) and Lemma 2.20 of Lauritzen (1996).

To show part (d), suppose that  $C^*$  is a clique of  $g'$ . Then  $C^*$  is complete in  $g$ , so  $C^* \subset C_l$  for some  $l \in \{1, \dots, k\}$ . If  $C^* \subset C_q$  then part (b) implies that either  $i \notin S_q$  or  $j \notin S_q$ . So either  $C^* \subset C_{q_1}$  or  $C^* \subset C_{q_2}$ . Hence  $C^*$  is contained in at least one of  $C_1, \dots, C_{q-1}, C_{q_1}, C_{q_2}, C_{q+1}, \dots, C_k$ . Part (c) shows that  $C_1, \dots, C_{q-1}, C_{q_1}, C_{q_2}, C_{q+1}, \dots, C_k$  are complete sets in  $g'$  and the result follows.

The next lemma uses (4) and Theorem 6 to simplify (15).

**LEMMA 7.** *Suppose that  $g$  and  $g'$  are the decomposable graphs defined above. Then, using the notation of (12), and Theorem 6*

$$\begin{aligned}
& \frac{h(g, \delta, \Phi) h(g', \delta^*, \Phi^*)}{h(g', \delta, \Phi) h(g, \delta^*, \Phi^*)} \\
&= \frac{\left| \Phi_{DD|S_{q_2}} \right|^{\left( \frac{\delta + |S_{q_2}| + 1}{2} \right)} \left| \Phi_{ii|S_{q_2}}^* \right|^{\left( \frac{\delta^* + |S_{q_2}|}{2} \right)} \left| \Phi_{jj|S_{q_2}}^* \right|^{\left( \frac{\delta^* + |S_{q_2}|}{2} \right)}}{\left| \Phi_{ii|S_{q_2}} \right|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} \left| \Phi_{jj|S_{q_2}} \right|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} \left| \Phi_{DD|S_{q_2}}^* \right|^{\left( \frac{\delta^* + |S_{q_2}| + 1}{2} \right)}} \times \\
& \frac{\Gamma \left( \frac{\delta + |S_{q_2}|}{2} \right) \Gamma \left( \frac{\delta^* + |S_{q_2}| + 1}{2} \right)}{\Gamma \left( \frac{\delta + |S_{q_2}| + 1}{2} \right) \Gamma \left( \frac{\delta^* + |S_{q_2}|}{2} \right)}, \tag{23}
\end{aligned}$$

where  $D = \{i, j\}$ ,  $\Phi_{DD|S_{q_2}} = \Phi_{DD} - \Phi_{DS_{q_2}} (\Phi_{S_{q_2}S_{q_2}})^{-1} \Phi_{S_{q_2}D}$ , and  $\Phi_{ii|S_{q_2}}, \Phi_{jj|S_{q_2}}, \Phi_{DD|S_{q_2}}^*$ ,  $\Phi_{ii|S_{q_2}}^*$  and  $\Phi_{jj|S_{q_2}}^*$  are defined similarly.

PROOF. To obtain an expression for  $h(g', \delta, \Phi)$  we require the following technical lemma based on Lemma 2.13 of Lauritzen (1996).

LEMMA 8. Let  $\tilde{C}_1, \dots, \tilde{C}_k$  be a perfect sequence with separators  $\tilde{S}_2, \dots, \tilde{S}_k$ . Assume that  $\tilde{C}_t \subset \tilde{C}_p$  for some  $t \neq p$  and that  $p$  is minimal with this property for fixed  $t$ . Then

- (a) If  $p < t$  then  $\tilde{S}_t = \tilde{C}_t$  and  $\tilde{C}_1, \dots, \tilde{C}_{t-1}, \tilde{C}_{t+1}, \dots, \tilde{C}_k$  is a perfect sequence with separators  $\tilde{S}_2, \dots, \tilde{S}_{t-1}, \tilde{S}_{t+1}, \dots, \tilde{S}_k$   
 (b) If  $p > t$  then  $\tilde{S}_p = \tilde{C}_t$  and  $\tilde{C}_1, \dots, \tilde{C}_{t-1}, \tilde{C}_p, \tilde{C}_{t+1}, \dots, \tilde{C}_{p-1}, \tilde{C}_{p+1}, \tilde{C}_k$  is a perfect sequence with separators  $\tilde{S}_2, \dots, \tilde{S}_{t-1}, \tilde{S}_t, \tilde{S}_{t+1}, \dots, \tilde{S}_{p-1}, \tilde{S}_{p+1}, \tilde{S}_k$

**Proof** of Lemma 8. See Lemma 2.13 of Lauritzen and its proof.

From Lemma 8, a perfect sequence of complete sets  $\tilde{C}_1, \dots, \tilde{C}_k$  containing the cliques of  $g'$  can be thinned by removing complete sets that are not cliques and reordering the sequence. From Lemma 8, the right-hand side of (4) is invariant to this thinning process. Successive application of the thinning process gives a perfect sequence consisting of the cliques of  $g'$ .

From (4), Theorem 6 and Lemma 8

$$h(g', \delta, \Phi) = \frac{\prod_{i=1, \dots, q-1, q_1, q_2, q+1, \dots, k} \left[ \left| \frac{\Phi_{C_i C_i}}{2} \right|^{\left( \frac{\delta + |C_i| - 1}{2} \right)} \Gamma_{|C_i|} \left( \frac{\delta + |C_i| - 1}{2} \right)^{-1} \right]}{\prod_{i=2, \dots, q-1, q_1, q_2, q+1, \dots, k} \left[ \left| \frac{\Phi_{S_i S_i}}{2} \right|^{\left( \frac{\delta + |S_i| - 1}{2} \right)} \Gamma_{|S_i|} \left( \frac{\delta + |S_i| - 1}{2} \right)^{-1} \right]}. \quad (24)$$

Now consider the ratio  $h(g, \delta, \Phi)/h(g', \delta, \Phi)$ . Simplifying the expressions from (4) and (24) gives

$$\frac{h(g, \delta, \Phi)}{h(g', \delta, \Phi)} = \frac{|\Phi_{C_q C_q}|^{\left( \frac{\delta + |S_{q_2}| + 1}{2} \right)} |\Phi_{S_q S_q}|^{\left( \frac{\delta + |S_{q_2}| - 1}{2} \right)} \Gamma \left( \frac{\delta + |S_{q_2}|}{2} \right)}{|\Phi_{C_{q_1} C_{q_1}}|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} |\Phi_{C_{q_2} C_{q_2}}|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} \Gamma \left( \frac{\delta + |S_{q_2}| + 1}{2} \right) 2\sqrt{\pi}}. \quad (25)$$

Substituting

$$\begin{aligned} |\Phi_{C_q C_q}| &= |\Phi_{DD|S_{q_2}}| |\Phi_{S_{q_2}}| \\ |\Phi_{C_{q_1} C_{q_1}}| &= |\Phi_{ii|S_{q_2}}| |\Phi_{S_{q_2}}| \\ |\Phi_{C_{q_2} C_{q_2}}| &= |\Phi_{ii|S_{q_2}}| |\Phi_{S_{q_2}}| \end{aligned}$$

into (25) gives

$$\frac{h(g, \delta, \Phi)}{h(g', \delta, \Phi)} = \frac{|\Phi_{DD|S_{q_2}}|^{\left( \frac{\delta + |S_{q_2}| + 1}{2} \right)} \Gamma \left( \frac{\delta + |S_{q_2}|}{2} \right)}{|\Phi_{ii|S_{q_2}}|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} |\Phi_{jj|S_{q_2}}|^{\left( \frac{\delta + |S_{q_2}|}{2} \right)} \Gamma \left( \frac{\delta + |S_{q_2}| + 1}{2} \right) 2\sqrt{\pi}}.$$



A similar expression can be derived for the ratio  $h(g, \delta^*, \Phi^*)/h(g', \delta^*, \Phi^*)$  and the result follows.

The following lemma gives an efficient method for evaluating the terms in (23) using Cholesky decompositions.

LEMMA 9. *Using the notation of Theorem 6 and Lemma 7, suppose that the matrix  $A_{C_q C_q} > 0$  is partitioned as*

$$A_{C_q C_q} = \begin{pmatrix} A_{S_{q_2} S_{q_2}} & A_{S_{q_2} D} \\ A_{D S_{q_2}} & A_{DD} \end{pmatrix}$$

and has Cholesky decomposition  $A_{C_q C_q} = LL'$  where

$$L = \begin{pmatrix} L_{S_{q_2} S_{q_2}} & 0 \\ L_{D S_{q_2}} & L_{DD} \end{pmatrix}$$

and

$$L_{DD} = \begin{pmatrix} l_{\alpha\alpha} & 0 \\ l_{\beta\alpha} & l_{\beta\beta} \end{pmatrix}.$$

Then

- (a)  $A_{DD|S_{q_2}} = L_{DD} (L_{DD})'$
- (b)  $|A_{DD|S_{q_2}}| = (l_{\alpha\alpha})^2 (l_{\beta\beta})^2$
- (c)  $A_{\alpha\alpha|S_{q_2}} = (l_{\alpha\alpha})^2$
- (d)  $A_{\beta\beta|S_{q_2}} = (l_{\beta\alpha})^2 + (l_{\beta\beta})^2$

PROOF. The proof is straightforward and is omitted.

Equation (15) and parts (b)—(d) of Lemma 9 give an efficient expression for the conditional distributions in Section 4. The main computational effort is in updating the Cholesky decompositions of the matrices  $\Phi_{C_q C_q}$  and  $\Phi_{C_q C_q}^*$  whenever an edge is added or deleted. From Lemma 9, these Cholesky decomposition must be done with the entries for the  $i$ th and  $j$ th vertices in the lower right corner. Note that efficient Cholesky updating routines using Givens rotations are available in Matlab and Fortran. Note also that the dimensions of  $\Phi_{C_q C_q}$  and  $\Phi_{C_q C_q}^*$  depend on the cliques sizes and may be much smaller than  $p$ . Thus our method has the local computational properties described in Giudici and Green (1999) and will have similar computational cost to their method per iteration of the Gibbs sampler.

## Acknowledgement

The research of Robert Kohn and Helen Armstrong was partially supported by an Australian Research Council Grant.

## References

Atay-Kayis, A. and H. Massam (2005). A Monte Carlo method to compute the marginal likelihood in non decomposable graphical gaussian models. *Biometrika* in press.

- Barnard, J., R. McCulloch, and X. Meng (2000). Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* 10, 1281–1311.
- Brooks, S., P. Giudici, and G. O. Roberts (2003). Efficient construction of reversible jump Markov chain Monte Carlo proposal distributions (with discussion). *J. Royal Statistical Society B* 65(1), 3–55.
- Castelo, R. and N. Wormald (2001). Enumeration of  $p_4$ -free chordal graphs. *Journal of Graphs and Combinatorics* (in press), or Universiteit Utrecht, Technical Report UU-CS-2001-12, June 2001.
- Chiu, T., T. Leonard, and K. Tsui (1996). The matrix-logarithm covariance model. *Journal of the American Statistical Association* 81, 310–20.
- Dawid, A. (1981). Some matrix-variate distribution theory: notational considerations and a Bayesian application. *Biometrika* 68(1), 265–274.
- Dawid, A. P. and S. Lauritzen (1993). Hyper Markov laws in the statistical analysis of decomposable graphical models. *The Annals of Statistics* 21(3), 1272–1317.
- Dellaportas, P. and J. Forster (1999). Markov chain Monte Carlo model determination for hierarchical and graphical log-linear models. *Biometrika* 86(3), 615–633.
- Dellaportas, P., P. Giudici, and G. Roberts (2004). Bayesian inference for non-decomposable graphical Gaussian models. *Sankhya*, Series A.
- Dempster, A. (1972). Covariance selection. *Biometrics* 28, 157–175.
- Dempster, A. P. (1969). *Elements of Continuous Multivariate Analysis*. Reading, MA: Addison-Wesley.
- Drton, M. and M. D. Perlman (2004). Model selection for Gaussian concentration graphs. *Biometrika* 91(3), 591–602.
- Efron, B. and C. Morris (1976). Multivariate Empirical Bayes estimation of covariance matrices. *Annals of Statistics* 4, 22–32.
- Frydenberg, M. and S. Lauritzen (1989). Decomposition of maximum likelihood in mixed interaction models. *Biometrika* 76(3), 539–555.
- Geiger, D. and D. Heckerman (2002). Parameter priors for directed acyclic graphical models and the characterization of several probability distributions. *Annals of statistics* 30(5), 1412–1440.
- Giudici, P. (1996). Learning in graphical Gaussian models. In A. P. D. J. Berger, J. M. Bernardo and A. F. M. Smith (Eds.), *Bayesian Statistics 5: Proceedings of the Fifth Valencia International Meeting, June 5-9, 1994*, pp. 621–628. Oxford University Press.
- Giudici, P. and R. Castelo (2003). Improving Markov chain Monte Carlo model search for data mining. *Machine learning* 50, 127–158.
- Giudici, P. and P. J. Green (1999). Decomposable graphical Gaussian model determination. *Biometrika* 86(4), 785–801.

- Jones, B., C. Carvalho, A. Dobra, C. Hans, C. Carter, and M. West (2005). Experiments in stochastic computation for high-dimensional graphical models. Preprint.
- Lauritzen, S. L. (1996). *Graphical models*. Oxford University Press.
- Liechty, J. C., M. W. Liechty, and P. Müller (2004). Bayesian correlation estimation. *Biometrika* 91(1), 1–14.
- Mardia, K. V., J. T. Kent, and J. M. Bibby (1979). *Multivariate Analysis*. London: Academic Press.
- Muirhead, R. (1982). *Aspects of Multivariate Statistical Theory*. Wiley.
- Roverato, A. (2000). Cholesky decomposition of a hyper inverse Wishart matrix. *Biometrika* 87, 99–112.
- Roverato, A. (2002). Hyper inverse Wishart distribution for non-decomposable graphs and its application to Bayesian inference for Gaussian graphical models. *Scandinavian Journal of Statistics* 29, 391–411.
- Smith, M. and R. Kohn (2002). Bayesian parsimonious covariance matrix estimation for longitudinal data. *Journal of the American Statistical Association* 87, 1141–1153.
- Wong, F., C. Carter, and R. Kohn (2003). Efficient estimation of covariance selection models. *Biometrika* 90, 809–830.
- Wormald, N. (1985). Counting labelled chordal graphs. *Graphs and combinatorics* 1, 193–200.
- Yang, R. and J. Berger (1994). Estimation of a covariance matrix using the reference prior. *Annals of Statistics* 22, 1195–1211.