Bayesian Estimation for Gaussian Graphical Models: Structure Learning, Predictability, and Network Comparisons

Donald R. Williams

University of California, Davis

Gaussian graphical models (GGM; "networks") allow for estimating conditional independence structures that are encoded by partial correlations. This is accomplished by identifying non-zero relations in the inverse of the covariance matrix. In psychology the default estimation method uses ℓ_1-regularization, where the accompanying inferences are restricted to frequentist objectives. Bayesian methods remain relatively uncommon in practice and methodological literatures. To date, they have not yet been used for estimation and inference in the psychological network literature. In this work, I introduce Bayesian methodology that is specifically designed for the most common psychological applications. The graphical structure is determined with posterior probabilities, which allow for assessing conditional dependent and independent relations. Additional methods are provided for extending inference to specific aspects within- and between-networks, including partial correlation differences and Bayesian methodology to quantify network predictability. I first demonstrate that the decision rule based on posterior probabilities can be calibrated to the desired level of specificity. The proposed techniques are then demonstrated in several illustrative examples. The methods have been implemented in the R package BGGM.

Keywords: Gaussian Graphical Model, Bayesian, Partial Correlation, Prediction, Explanation, Bayesian R^2, Networks

Introduction

Gaussian graphical models (GGM) have become increasingly popular in the psychological literature (Epskamp, Kruis, & Marsman, 2017; Epskamp, Waldorp, Mottus, & Borsboom, 2018). In contrast to structural equations models, an ubiquitous technique in the behavioral sciences, GGMs do not characterize hypothesized relations between latent factors and observed variables (Borsboom & Cramer, 2013). Rather, all variables are allowed to co-vary with relations encoded in the off-diagonal elements of the precision matrix (i.e., the inverse of the covariance matrix; Dempster, 1972; Whittaker, 1990). When the covariances are standardized and the sign ($\pm$) reversed, assuming multivariate normality, this results in partial correlations (ρ_{ij}) that are conditionally dependent relationships ($\rho_{ij} \neq 0$) between two variables, wherein the effects of all other variables have been controlled for (Baba, Shibata, & Sibuya, 2004; Baba & Sibuya, 2005). For example, in the context of a clinical application, GGMs allow for characterizing direct effects between a potentially arbitrary number of symptoms (Cao et al., 2018; Hajdúk, Klein, Harvey, Penn, & Pinkham, 2018). Because direct effects are suggestive of causal pathways (Pearl, 2009), this has contributed to the increased interest in GGMs. This emerging line of psychological inquiry is often referred to as "partial correlation networks," and they have been used to characterize several constructs that include political attitudes (Dalege, Borsboom, van Harreveld, & van der Maas, 2017), psychosis (Isvoranu et al., 2017; van Rooijen et al., 2017), post-traumatic stress disorder (Armour, Fried, Deserno, Tsai, & Pietrzak, 2017; McNally et al., 2015), substance abuse (Rhemtulla et al., 2016), and well-being (Deserno, Borsboom, Begeer, & Geurts, 2017).

Moreover, inquires have not been restricted to psychological applications, but there has also been substantial research into estimation methods that extends well beyond psychology (Fan, Liao, & Liu, 2016; Kuismin & Sillanpää, 2017). To my knowledge, most novel statistical methods for GGMs are developed to address a very specific problem that is not common in psychological applications, wherein there are more variables (p) than observations (n). The *so-called* large p small n problem that is commonplace when studying regulatory networks (Kuismin & Sillanpää, 2017), for example

microarray expression profiles for potentially thousands of genes (Mao, Van Hemert, Dash, & Dickerson, 2009; Wille et al., 2004), or when applying functional magnetic resonance imaging to capture the dimensionality of the whole-brain (Das et al., 2017; Luo, 2014). In these situations some form of regularization is *necessary*, because the covariance matrix cannot be inverted ($n < p$), in addition to customary estimators (i.e., maximum likelihood) becoming unstable when approaching high-dimensional settings ($p \rightarrow n$; Khondker, Zhu, Chu, Lin, & Ibrahim, 2013).

Despite regularization not being *necessary*, the default approach for estimating psychologyical "networks" is the graphical lasso ("least absolute shrinkage and selection operator"; glasso) which makes use of ℓ_1-regularization (R. Tibshirani, 1996). In the familiar setting of regression, this technique adds a penalty to the customary sum of squares, whereas the penalty is applied to the covariances in multivariate settings (Friedman, Hastie, & Tibshirani, 2008). While this approach overcomes the large p small n problem, it has primarily gained popularity in psychology because it pushes smaller estimates to exactly zero, thereby achieving an important theoretical goal of estimating a sparse network (Epskamp & Fried, 2016).

There are several important limitations of lasso (and related methods), especially as it relates to its use in applied settings. It should first be noted that ℓ_1-regularization is a relatively new technique, and thus remains an active area of research in the statistical literature (Avagyan, Alonso, & Nogales, 2017; Cai, Liu, & Luo, 2011; Mazumder & Hastie, 2012). While the ℓ_1-penalty can indeed push parameters to zero, strong assumptions are *necessary* for model selection consistency (Zhao & Yu, 2006). That is, even with infinite data ($n \rightarrow \infty$), if *only* minimizing the sums of squares with respect to the ℓ_1-penalty, the probability of selecting the true model will not often tend to 1 (Williams & Rast, 2018; Zhao & Yu, 2006). This stands in contrast to traditional approaches in psychology that are known to be consistent, such nonreguluarized model selection with the Bayesian information criteria (Burnham & Anderson, 2004; Casella, Girón, Martinez, & Moreno, 2009). Indeed, in data structures more common to psychology, it was shown that ℓ_1-regularization has distinct limitations compared to nonregularized estimation (see: Williams & Rast, 2018; Williams, Rhemtulla, Wysocki, & Rast, 2019).

Perhaps the most notable limitation, when considering that psychology as a field is explicatory in nature (Yarkoni & Westfall, 2017), is that ℓ_1-regularization presents challenges for statistical inference (Hastie, Tibshirani, & Wainwright, 2015). In this context, since lasso is a frequentist approach, by valid I am referring to calibrated long-run frequencies (e.g., α; Neyman, 1935). Of note, ℓ_1-regularization only provides point estimates from convex optimization (i.e., no significance testing is involved). Most research has fo-

cused exclusively on point estimation and rates of convergence, with only recently p-values (Lockhart, Taylor, Tibshirani, & Tibshirani, 2014) and confidence intervals (CI) being developed (Janková & van de Geer, 2015, 2017). These approaches range from post-selection adjusted p-values (Taylor & Tibshirani, 2017; R. J. Tibshirani, Taylor, Lockhart, & Tibshirani, 2016) to obtaining valid measures of uncertainty (Zhang & Zhang, 2014), which are not readily available with customary bootstrap schemes (Bühlmann, Kalisch, & Meier, 2014). A recent approach for constructing CIs, which allows for claiming there is a significant effect, is the desparsified ℓ_1-penalized estimator (Van De Geer, Bühlmann, Ritov, & Dezeure, 2014). In other words the theoretical motivation (estimating exact zeroes), in psychology, for using ℓ_1-regularization is non-existent when attempting to make inference with null hypothesis significance testing. They are also primarily characterized in high-dimensional settings (Janková & van de Geer, 2017), which suggests that not only is generalizability of concern, but it is unclear whether they have advantages compared to nonregularized methods (Williams & Rast, 2018; Williams, Rhemtulla, et al., 2019).

These limitations partially motivate this work, in that I introduce a (nonregularized) Bayesian approach that should converge on the true model with minimal assumptions, does not require manual adjustments *after* model selection to reduce the false positive rate, and readily provides a valid measure of uncertainty. Further, an additional advantage is that Bayesian inference focuses on posterior probabilities (Section: Decision Rule), in addition to providing posterior distributions for each partial correlation. This allows for probabilistically learning of conditional (in)dependence structures, and straightforward extensions specifically developed for the most common psychological applications ($p < n$).

Recently, the popularity of Bayesian methods has grown in certain psychological applications (Wagenmakers, Love, et al., 2018; Wagenmakers, Marsman, et al., 2018). To our knowledge, however, Bayesian methods have only been used to estimate partial correlations between two variables in the psychological literature (Nuijten, Wetzels, Matzke, Dolan, & Wagenmakers, 2015; Wetzels & Wagenmakers, 2012). The only related example used Bayesian methods for estimating a directed acyclic graph (McNally, Heeren, & Robinaugh, 2017), whereas GGMs encode *undirected* relationships between variables. This relative absence of Bayesian methods is not restricted to psychology; for example, there are only two R packages for estimating GGMs. Neither provides full posterior distributions for the partial correlations (Leday & Richardson, 2018; Mohammadi & Wit, 2015). This limits the capabilities for psychological applications, in that the posterior distributions allow for computing additional information such as measures of uncertainty for variance explained and predictive accuracy for each variable in the network (Section: Network Predictability). This is particularly important,

because there is growing interest in psychology to consider the predictive capabilities of our models (Yarkoni & Westfall, 2017). By considering the predictive accuracy of individual variables in the network, for example, this may provide more targeted interventions for clinical applications. Indeed, there is an emerging interest in evaluating predictive accuracy of specific variables in psychological networks (Haslbeck & Fried, 2017; Haslbeck & Waldorp, 2018). However, variable importance is often inferred from point estimates which in my view can lead to overconfident inferences (Birkeland, Blix, Solberg, & Heir, 2017). Together, the present work is not only a valuable contribution to the GGM literature in general, but also provides novel Bayesian methodology specifically designed for psychological applications. These include measures of network predictability, with respect to the posterior distributions, as well as learning the conditional (in)dependence structure of psychological constructs.

The aim of the present work is to introduce Bayesian methodology for estimating Gaussian graphical models, in addition to describing techniques that can be used for explanation and prediction. I first introduce the customary notation for GGMs that will be used for the remainder of this work, and then describe the rationale for using nonregularized estimation. The proposed methodology is then outlined, which makes use of the Wishart prior distribution that is conjugate to the precision matrix, as well as the decision rule for determining conditional relationships. I then describe two techniques in which the estimated precision matrix can be used to compute partial correlation differences, variance explained, and out-of-sample predictive accuracy for each variable in the network. In the following sections, the techniques are employed in a series of illustrative examples. I end by discussing future directions as well as limitations of this work.

The Gaussian graphical model

The Gaussian graphical model captures undirected, conditional relationships (Lauritzen, 1996), that are typically visualized to infer the underlying structure (i.e., the "network"; Højsgaard, Edwards, & Lauritzen, 2012). The undirected graph is $G = (V, E)$, and includes a vertex set $V = \{1, ..., p\}$ as well as an edge set $E \subset V \times V$. Let $\mathbf{Y} = (\mathbf{y}_1, ..., \mathbf{y}_p)^\top$ be a $n \times p$ matrix, where each $\mathbf{y}$ is a n dimensional vector that is indexed by the graphs vertices. Assume they follow a multivariate normal distribution $\mathbf{Y} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with the mean vector $\boldsymbol{\mu} = (0_1, ..., 0_p)^\top$ and $p \times p$ positive definite covariance matrix $\boldsymbol{\Sigma}$. Denote the precision matrix $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$. The graph is obtained from the off-diagonal elements $\theta_{ij} \in \boldsymbol{\Theta}_{ij}$. This is used to construct an adjacency matrix A that follows

$$A_{ij} = \begin{cases} 1, & \text{if } \theta_{ij} \neq 0, \ 1 \leq i < j \leq p \\ 0, & \text{otherwise}, \end{cases} \tag{1}$$

with $1 \leq i < j \leq p$ denoting the elements in the upper-triangular of the $p \times p$ matrix. Further, $(i, j) \in E$ when the variables i and j are *not* conditionally independent and set to zero otherwise. Note that the selected edges, when standardized, are partial correlations (ρ) determined to be non-zero. These are computed directly from the precision matrix–i.e.,

$$\rho_{ij} = \frac{-\theta_{ij}}{\sqrt{\theta_{ii}\theta_{jj}}}, \ 1 \leq i < j \leq p. \tag{2}$$

In this work, these partial correlations are explicitly used for the methods that draw samples from the posterior distribution, whereas the precision matrix is primarily used for the methods that have an analytic solution.

There are often many parameters estimated in psychological networks–i.e., the elements of $\boldsymbol{\Theta}$. This does raise the possibility of overfitting, which can be assuaged with regularization (McNeish, 2015). However, it should be noted that estimating $\boldsymbol{\Theta}$ is straightforward in low-dimensional settings (Williams & Rast, 2018; Williams, Rhemtulla, Wysocki, & Rast, 2018). However, even in low dimensional settings ($p < n$) estimation error increases when p approaches n (e.g., mean squared error). In other words, n must be sufficiently larger than p for the $\boldsymbol{\Theta}$ to be estimated accurately. As shown in Williams et al. (2018), glasso does not show clear benefits (even for predictive accuracy) compared to the maximum likelihood estimate with dimensions representative of psychological networks (e.g., $p = 20$, $n = 250$; Hajdúk et al., 2018; McNally et al., 2017, 2015). Together, while regularized networks have emerged as the default approach in psychology (Epskamp & Fried, 2016), the method described below is nonregularized by default.

Bayesian Estimation of Gaussian Graphical Models

Bayesian methods often lack scalability to complex models with numerous parameters to be estimated (Barrientos & Peña, 2017; Newton, Polson, & Xu, 2018). This is due to computational demand, which can make sampling from the posterior distributions not only slow but also inefficient (Chopin, Gadat, Guedj, Guyader, & Vernet, 2015). This is important to consider, in psychological applications, because with even 20 variables there are 190 partial correlations in total. To overcome this issue, conjugate prior distributions can be utilized, where the posterior and prior are from the same family of probability distributions (Gutiérrez-Peña et al., 1997). This is advantageous in that sampling from the posteriors is expedited, thereby allowing for scaling to larger dimension settings.

A popular Bayesian method for GGMs makes use of the G-Wishart distribution, which is conjugate to $\boldsymbol{\Theta}$. This is a generalization of the Wishart distribution, where the off-diagonal elements of $\boldsymbol{\Theta}$ can be constrained to zero (Atay-Kayis & Massam, 2005; Dobra, Lenkoski, & Rodriguez,

2010). Some of the earliest approaches employed reversible jump MCMC (Green, 1995), for example in Wang (2012) and Lenkoski (2013), whereas the most recent sampling strategies use a birth and death MCMC algorithm (Mohammadi & Wit, 2015). This provides the customary framework for Bayesian model selection, in that covariance selection uses an indicator function analogous to spike and slab prior distributions that are commonly used for general linear models (George & McCulloch, 1993; O'Hara & Sillanpää, 2009; Rouder, Haaf, & Vandekerckhove, 2018). That is, the graph can be selected with either Bayesian model averaging (Hoeting, Madigan, Raftery, & Volinsky, 1999) or maximum a posterior probability (Barbieri & Berger, 2004). Importantly, the software implementations of the *G*-Wishart approach do not return full posteriors for the parameters, but instead they provide only point estimates. For the present purposes, where our aim is to provide tools specifically for psychological applications, this limits extensions to ask additional questions related to the estimated network (e.g., edge difference and network predictability)

Wishart Prior Distribution

The proposed method uses the Wishart prior distribution that is directly conjugate for $\boldsymbol{\Theta}$ (Kubokawa & Srivastava, 2008). In Bayesian mixed-effects models, which are commonly used in psychology, the inverse-Wishart distribution is often used to estimate the covariance matrix $\Sigma = \boldsymbol{\Theta}^{-1}$ (Schuurman, Grasman, & Hamaker, 2016). The present application stands in contrast to more common methodological inquiries into the properties of the inverse-Wishart distribution. It is common to examine 2×2 matrices (Alvarez, Niemi, & Simpson, 2014; Liu, Zhang, & Grimm, 2016), for example the covariance between random slopes and intercepts, whereas the proposed method is for much larger dimensions. Nonetheless, as described below, the Wishart distribution is not only efficient but also provides an accurate estimate for $\boldsymbol{\Theta}$ (Section: Decision Rule).

The joint posterior density for the precision matrix follows

$$p(\boldsymbol{\Theta}|\mathbf{Y}) \propto p(\mathbf{Y}|\boldsymbol{\Theta})p(\boldsymbol{\Theta}), \tag{3}$$

where $\mathbf{Y}$ is a $n \times p$ matrix drawn from a multivariate normal distribution–i.e.,

$$\mathbf{Y} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Theta}^{-1}). \tag{4}$$

When using the conjugate Wishart prior, $\mathcal{W}(k, \epsilon\mathbf{I}_p)$, with k degrees of freedom and identity matrix $\mathbf{I}_p$, the posterior distribution also has a Wishart distribution–i.e.,

$$\boldsymbol{\Theta}|\mathbf{Y} \sim \mathcal{W}(k + n, (\boldsymbol{S} + \epsilon\mathbf{I}_p)^{-1}), \tag{5}$$

where $\boldsymbol{S}$ is the sums of squares matrix $\mathbf{Y}'\mathbf{Y}$ and ϵ is a constant. The posterior mode then has a closed form

$$\arg\max_{\boldsymbol{\Theta}} p(\boldsymbol{\Theta}|\mathbf{Y}) = (k + n - p - 1)(\boldsymbol{S} + \epsilon\mathbf{I}_p)^{-1}, \tag{6}$$

as well as the posterior posterior variance

$$\text{Var}(\boldsymbol{\Theta}|\mathbf{Y}) = (k + n)((\boldsymbol{S} + \epsilon\mathbf{I}_p)^{-1^2} + \boldsymbol{d}\boldsymbol{d}'), \tag{7}$$

where $\boldsymbol{d} = \text{diag}(\boldsymbol{S}+\epsilon\mathbf{I}_p)$. By setting $k = p+1$, and ϵ to a small value, say, 1.0×10^{-10}, results (approximately) in the maximum likelihood estimate $n(\boldsymbol{S}^{-1})$. Consequently, by assuming $p(\theta_{ij}|\mathbf{Y})$ is normally distributed (see pg. 88: "Asymptotic normality and consistency"; Gelman et al., 2014), allows for constructing (equal tailed) credible intervals and computing posterior probabilities. For example, the former can be obtained as

$$\int_l^u p(\theta_{ij}|\mathbf{Y})d\theta_{ij} = 1 - \alpha, \tag{8}$$

where l and u denote the lower and upper bounds of the interval. What remains is the choice of α, which as shown below, can be (approximately) calibrated to the desired level of specificity[1]. The latter follows

$$\Pr(\theta_{ij} > 0|\mathbf{Y}), \tag{9}$$

which corresponds to the posterior probability of a positive effect. Both can be computed analytically with the point estimate in (6) and variance in (7). This simple solution can be used for determining E and analytically deriving "network" predictability.

There is one drawback of this analytic form. Because there are no posterior samples, this limits its applicability for extending inference beyond the graphical structure.. Indeed, capturing uncertainty is a key aspect of this work compared to, say, ℓ_1-regularization and alternative Bayesian methods that also provide point estimates for $\boldsymbol{\Theta}$. I thus describe an additional approach for conveniently drawing samples from the posterior distribution. I implemented the noninformative Jeffreys' prior $|\boldsymbol{\Theta}|^{(p+1)/2}$. This allows for posterior sampling directly from a Wishart distribution–i.e.,

$$\boldsymbol{\Theta}|\mathbf{Y} \sim \mathcal{W}(n - 1, \boldsymbol{S}^{-1}). \tag{10}$$

---

[1]Specificity is a diagnostic measure commonly used to assess the performance of methods used to estimate GGMs. It is equal to $1 - \alpha$.

Table 1

Correspondence between directional posterior probabilities, α, and specificity.

$p(\rho_{ij} > 0|\mathbf{Y})$	α	SPC
0.9975	0.005	0.995
0.995	0.01	0.99
0.975	0.05	0.95
0.95	0.10	0.90

These samples, $s = 1, ..., S$, are used to compute the posterior distribution for the $p \times p$ partial correlation matrix, with $\rho_{ij} \in \mathcal{P}$–i.e.,

$$\mathcal{P}^{(s)} = -([\text{diag}(\boldsymbol{\sigma})^{(s)}]^{-1}\boldsymbol{\Theta}^{(s)}[\text{diag}(\boldsymbol{\sigma})^{(s)}]^{-1}), \qquad (11)$$

where $\boldsymbol{\sigma}$ are the square roots of $\text{diag}(\boldsymbol{\Theta})$ (Barnard, McCulloch, & Meng, 2000) and multiplying by -1 reverses the direction ($\pm$) as in (2). Because ρ is a standardized effect, and thus each ρ_{ij} is on the same scale, this allows for defining a *region of practical equivalence* (Kruschke, 2011). Rather than determining E in reference to 0, as in (8) and (9), a neighborhood around zero is defined (a null area; Betancourt, 2018). Consequently, given some region, $|\rho_{ij}| \leq \vartheta_0$, there is support for null values when the posterior probability is above a pre-determined threshold,

$$\int_{-\vartheta_0}^{\vartheta_0} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} > 1 - \alpha, \qquad (12)$$

where $p(\rho_{ij}|\mathbf{Y})$ is the posterior distribution. Conversely,

$$\int_{-\infty}^{-\vartheta_0} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} + \int_{\vartheta_0}^{\infty} p(\rho_{ij}|\mathbf{Y})d\rho_{ij} > 1 - \alpha, \qquad (13)$$

denotes the posterior probability *outside* of the region of practical equivalence exceeds $1 - \alpha$. This allows for both determining conditional dependent and practically independent relations, as well as comparing two edges (e.g., $\rho_{1,2} - \rho_{1,3}$).

Decision Rule

The Wishart distribution does not produce a sparse estimate for $\boldsymbol{\Theta}$. This stands in contrast to ℓ_1-regularization that provides exact zeroes. Thus, an explicit decision rule is required for estimating the graphical structure. I view this as advantageous, in that the evidence can be weighed with respect to the question at hand. To our knowledge, there have been two different decision rules applied to the MAP estimate (Equation: 6). The first, described in Kuismin and Sillanpää (2016), used the extended Bayesian information criteria (EBIC) for setting values to zero. This was accomplished with a step-wise procedure in which an algorithm sequentially removed edges until EBIC could not be further reduced. The second approach, specifically for high-dimensional data ($n \ll p$), introduced an analytic solution for Bayes factors in combination with a null sampling distribution for controlling the false discovery rate (Leday & Richardson, 2018). Both of these approaches are not strictly Bayesian, whereas the presented method obtains E from the estimated posterior probabilities that are conditional only on the data and the fitted model.

I determine E with directional posterior probabilities. This approach does not test a point null hypothesis ($\mathcal{H}_0 : \rho_{ij} = 0$) against an alternative hypothesis ($\mathcal{H}_1 : \rho_{ij}$). Rather, relations are inferred in reference to the posterior distribution on either side of zero (Figure 1; panel A). Adopting this perspective maps directly onto specificity which is a common metric for evaluating edge set identification. To see this, consider, $p(\rho_{ij} > 0|\mathbf{Y}) = 0.975$, which is the posterior probability the edge is greater than zero. It follows that $1 - 0.975$ is the posterior density less than zero. This is often referred to as the Bayesian p-value. The decision rule can then be calibrated to the desired level of specificity (SPC)–i.e.,

$$\begin{aligned} \text{SPC} &= \frac{\#TN}{\#TN + \#FP} \qquad (14) \\ &= 1 - \text{FPR}, \end{aligned}$$

where $\#TN$ and $\#FP$ are the number of true negatives and false positives, respectively, and FPR is the false positive rate (i.e., α). In other words, by using the relationship between α and $1 - p(\rho_{ij} > 0|\mathbf{Y})$ (for determining positive relations), it is possible to demonstrate that the proposed decision rule can be calibrated to SPC (at least approximately so).

I conducted a brief simulation to make this point clear. I assumed an identity matrix I_p ($p = 20$) for $\boldsymbol{\Theta}$ and increased the sample sizes $n \in \{100, 250, 500, 1000, 2500, 5000\}$. Because I_p is a null matrix, this allows for examining the global error rate for $\mathcal{G}$. Since there are 195 covariances $\frac{p(p-1)}{2}$ equal to zero, the number of false positives is approximately $\alpha \frac{p(p-1)}{2}$. For each sample size, the error rates were average across 50 simulation trials. The results for both the analytic solution (Equation 6) and the sampling based approach (Equation 10) are displayed in Figure 1. The desired level of specificity was set to 0.90, 0.95, and 0.99 ("Decision Rule"), which can be seen in panel B. The dotted lines correspond to the expected false positive rate (1 - SPC). The analytic solution was consistently at the nominal level, whereas the sampling based approach had a (slightly) higher error rate that ultimately converged to the desired level of specificity as n increased. The difference between approaches could be attributed to the shape of the respective distributions. That is, for the analytic solution, E is computed from the off-diagonal elements of $\boldsymbol{\Theta}$ that are normally distributed, whereas the sampling based approach uses the partial correlations that could be skewed. Importantly, as revealed in this simulation, both approaches can be (approximately) calibrated to the desired level of specificity.

Illustrative Example

In this section, I use the proposed methodology to estimate the underlying conditional (in)dependence structure of post-traumatic stress symptoms (i.e., PTSD; McNally et al., 2015). There are 20 variables (p) and 221 observations (n)
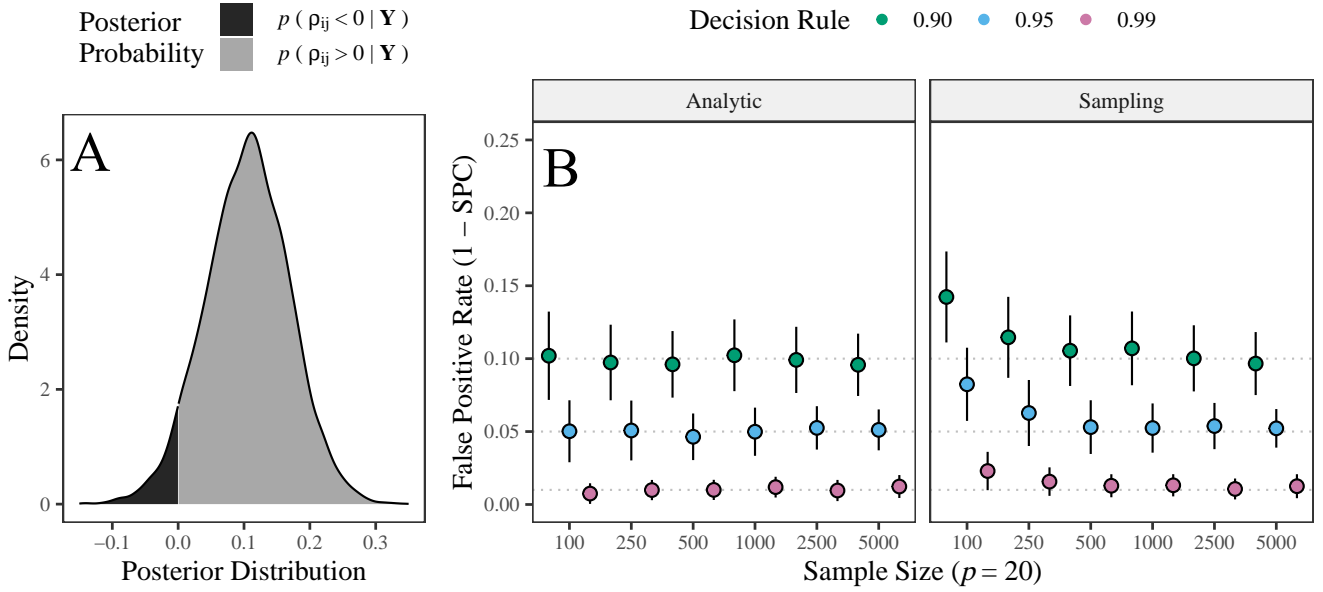
Figure 1. A) This panel shows an example posterior distribution for an edge (partial correlation). The shaded regions correspond to directional posterior probabilities. In this example, the probability of a positive effect was 95 % (thus the probability of a negative effect was 5 %). These probabilities are used for determining the network structure–a connection is included when the probability exceeds a threshold. B) This panel show that the decision rule based on probabilities can be calibrated to a desired level of specificity. The *y*-axis is the false positive rate which corresponds to 1 - specificity. The dotted line is the expect false positive rate. The analytic solution is provided in Equations (6 and 7). The sampling based approach is provided in Equations (10 and 11).

in total, each of which was measured on the 5 values ordinal scale. For demonstrative purposes, I also included the most popular estimation method in psychology–i.e., the graphical lasso with the extended Bayesian information criterion (glasso$_{EBIC}$ Epskamp, 2016). The ℓ_1-regularized networks were estimated with the R package **qgraph** using the default settings, in addition to **bootnet** for bootstrapping the estimated networks (Epskamp, Borsboom, & Fried, 2018). The Bayesian models were fitted with the package **BGGM** (Williams & Mulder, 2019b).

Structure Learning

Conditional Dependence. By structure learning I am referring to the edge set E or the conditional *dependence* structure of the data ($\rho_{ij} \neq 0$). For the Bayesian model three decision rules were used (Figure 1). These corresponded to expected specificity of 0.90, 0.95, and 0.99. Figure 2 (panel A) displays the estimated networks. There are clear differences between glasso$_{EBIC}$ and the Bayesian methods. The former estimated what could be considered a dense network, in which almost half of the nodes shared a connection (≈ 0.52 %). On the other hand, with the most conservative threshold for BGGM ("Post prob: 0.999"), the estimated structure had far (far) fewer connections than glasso$_{EBIC}$. Specifically, there was *only* 11 % connectivity. This discrepancy could be

due to the inflated false positive rate of ℓ_1-regularized estimation (Williams & Rast, 2018; Williams, Rhemtulla, et al., 2019).

Conditional Independence. A key aspect of Bayesian inference is the ability to assess the null hypothesis of conditional *independence* ($\rho_{ij} \neq 0$). In this case, based on the posterior distribution of the partial correlations, it is possible to compute the probability of null values (Equation 12). For demonstrative purposes, I followed Cohen's guidelines and defined a practically meaningful edge as what is customarily defined as a small effect size ($\rho = 0.1$ Cohen & Cohen, 1983). The decision rule was then set to 0.95–i.e., to be considered practically zero, 95 % of the probability density must lie within the region of practical equivalence ±0.1. Figure 2 (panel B) displays the estimated conditional dependence ("Practically Zero") and independence structures ("Practically Nonzero") estimated with the rope. For the latter, we see that there was actually no evidence for conditional independent relations in this data.[2] This stands in contrast to ℓ_1-regularized networks, where it might be tempting to conclude the zeros provide evidence for conditional

---

[2]I also varied the decision rule to consider various posterior probabilities. With a posterior probability of 0.85, there were a few conditionally independent effects. Of course, this also suggests a 15 % chance of a meaningful effect.
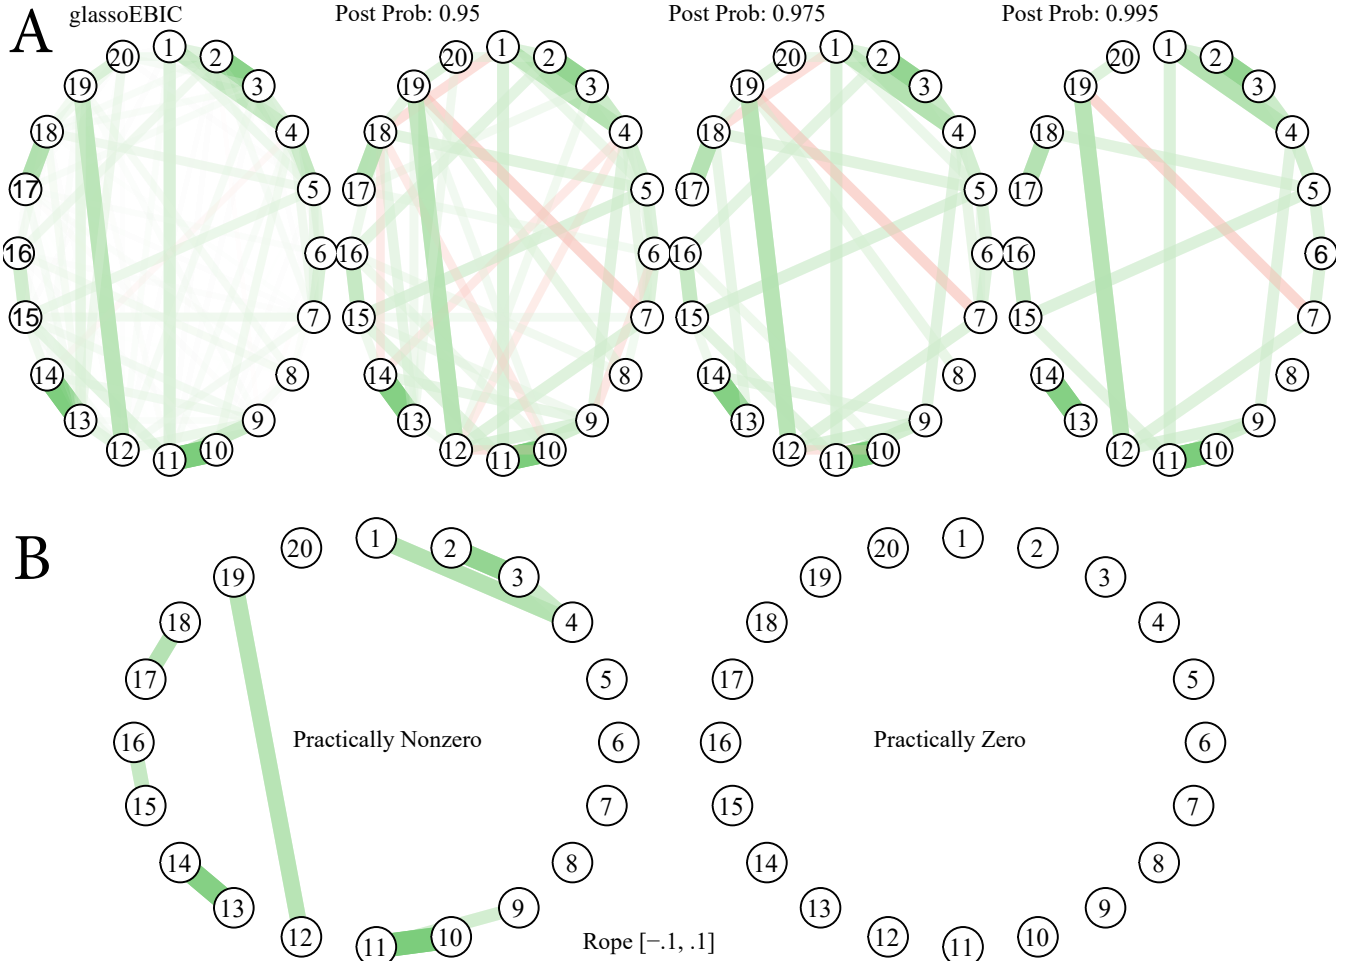
Figure 2. A) This panel shows the estimated networks with glasso$_{EBIC}$ and the Bayesian methods. For the latter, various posterior probabilities were used to determine the network structure. B) This panel only includes networks estimated with the proposed Bayesian methods. Rather than use a posterior probability in reference to zero, conditionally independent and dependent relations were determined with a region of practical equivalence (a null region). In this example, none of the edges were practically zero or conditionally independent. The node names are provided in Table 2.

independence. This is not much different than setting values to zero that were not statistically significant, which is commonly done in psychology and also does not provide evidence for the null hypothesis. The conditional dependence structure is also much sparser, which is perhaps not too surprising because a meaningful effect has been defined as 0.1.

Quantifying Uncertainty. As previously mentioned (Section: Introduction), ℓ_1-regularization provides point estimates for the edges. This makes extending inference beyond exploratory challenging. For non-regularized methods, as described in Williams et al. (2018), it is possible to bootstrap Θ. This is generally not advisable with ℓ_1-penalized estimates, in that a sparse solution has a point mass at zero that produces a distorted sampling distribution. However, in psychological applications, it is nonetheless common to conduct significance tests from bootstrapped ℓ_1-estimates. I thus

compared the Bayesian methods to bootstrapping glasso$_{EBIC}$. I generated 2,000 boostrap samples for glasso$_{EBIC}$ that took 18.6 minutes in total, whereas BGGM provided 2,000 posterior samples in less than 1 second. This difference arises because, for each bootstrap sample, glasso$_{EBIC}$ fits a separate model for each tuning parameter value (which controls the sparsity level) and then one is selected according to EBIC. As such, with 100 values (the default in **qgraph**), this results in 200,000 models being fitted to provide the same number samples as BGGM. For this dimension of data ($p = 20$ and $n = 221$), this demonstrates that sampling from the posterior can be accomplished very efficiently (Section: Wishart Prior Distribution).

The estimated edge weights are displayed in Figure 3 (Panel A), with the interval width set to 99 %. There are again clear difference between glasso$_{EBIC}$ and BGGM. For
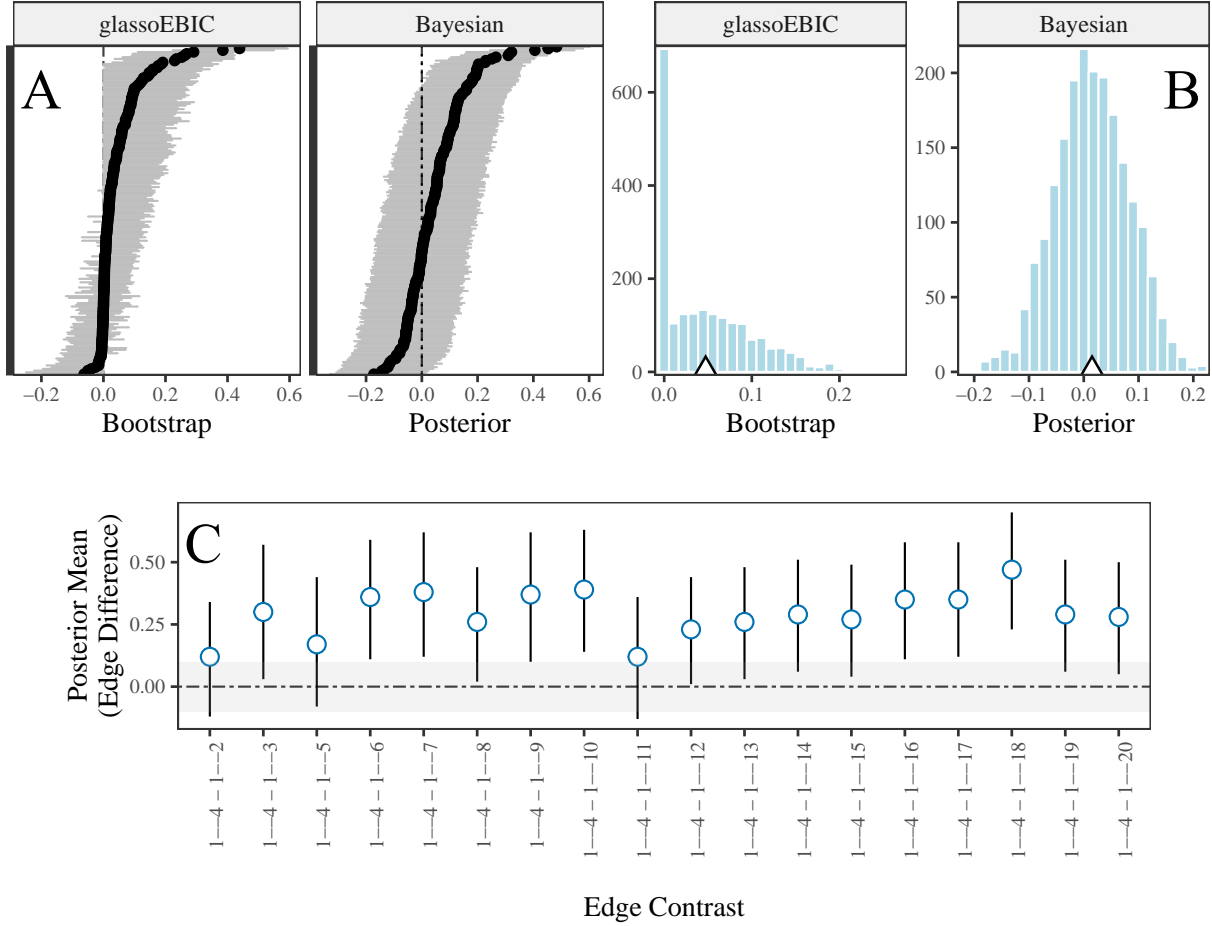
Figure 3. A) This panel includes the estimated edges from glasso$_{EBIC}$ and the Bayesian methods. The former uses a bootstrap approach, whereas the latter provides a posterior distribution. This is meant to highlight the differences in the shape of the distributions. That is, the regularized estimates are often bounded at zero and skewed (inferred from the location of the point estimate on the error bar). The error bars correspond to 99 % confidence and credible intervals. B) This panel further highlights the differences between the distributions. The ℓ_1-regularized distribution has a point mass at zero. This is a central challenge for constructing *valid* confidence intervals from ℓ_1-based estimates. I refer to Dezeure, Bühlmann, Meier, and Meinshausen (2014), and in particular the appendix, which shows how difficult it is to achieve correct coverage probabilities. On the other hand, the Bayesian methods can be calibrated to have nominal coverage levels and thus specificity (Figure 1). C) This panel includes edges differences for node 1 ("Intrusive Thoughts"). The contrasts are between the largest edge "1–4" compared to all the others. The grey area is the region of practical equivalence (±0.1).

example, it can been seen that the sampling distribution of the ℓ_1-penalized estimates depends on the partial correlation size. That is, if the value has sufficient separation from zero, the distribution was approximately symmetric. This was not the case for values closer to zero, where the distributions were highly skewed and had a lower boundary of zero. This can be inferred by noting where the mean is located in the distribution. To make this point clearer, I selected one edge from each model and displayed them Figure 2 (Panel B). Here the point mass at zero for glasso$_{EBIC}$ can readily be seen. In fact, while this edge was included in E, over 50 % of the boostrap samples were estimated as zero. On the

other hand, the posterior distributions were very similar and approximately normally distributed. Together, glasso$_{EBIC}$ could be ideal for determining the proportion of times an edge was selected, but based on the shape of the sampling distribution, it seems less than ideal for constructing intervals. I refer interested readers to Dezeure et al. (2014), and in particular the appendix, which shows how difficult it is to achieve correct coverage probabilities with ℓ_1-based estimates.

Edge Differences. It is important to extend inference beyond identifying the graphical structure and to, say, testing which edge has the most strength (i.e., the largest ef-

fect size). Edge differences have customarily been computed with **bootnet** from ℓ_1-regularized estimates. In this work, based on Equation 11, the posterior distributions for two edges can be subtracted from one another. This not only allows for determining whether the edges are different, but it is also possible to compute the posterior probability of practical equivalence. This is a powerful approach; for example, a researcher can ask whether certain edges are of the same strength for a given symptom. For demonstrative purposes, I tested whether edges were different or the same as the largest edge for node 1 or "Intrusive Thoughts" (Table 2). Here the largest relation was with node 4 or "Emotional Cue Reactivity."

Figure 3 includes the edge differences. They are summarized with the posterior means and 99 % credible intervals. Further, I also shaded the area from ±0.1 to illustrate the central idea behind the region of practical equivalence. That is, if the interval excludes the shaded region there is a meaningful difference (subjectively defined), whereas if the interval is completely within the shaded region, the two edges are practically equivalent. In this case, the edge "1–4" was larger than many but not all the others. In particular, for three edges the interval not only spanned +0.1 but they also included zero. These results highlight a key aspect of Bayesian inference, in that a measure of uncertainty is provided for the edges that allows for computing posterior probabilities of the difference or equivalence.

Network Predictability

The following is based on the correspondence between the elements of $\boldsymbol{\Theta}$ and multiple regression. Detailed proofs are provided in Stephens (1998), and further information can be found in Kwan (2014). In the context of GGMs, using regression to select edges is referred to as "neighborhood" selection (Meinshausen & Bühlmann, 2006). On the other hand, the method described here work directly with either the posterior distribution (Equation 10) or the maximum a posteriori (Equation 6). The technique can be described with traditional regression notation, in which p multiple regression models are fitted. First let each node V_j be defined as $\mathbf{y}$–i.e., the scores of the n subjects on the jth variable/node. Each node is regressed on the remaining $p-1$ variables, which estimates the potential "neighborhood" for each variable

$$\mathbf{y} = X\boldsymbol{\beta}_j + \boldsymbol{\varepsilon}, \tag{15}$$

where $\boldsymbol{\varepsilon}$ is an n-dimensional vector, with the mean as a vector of zeroes, and the covariance matrix as $\sigma^2\mathbf{I}_n$. Here X is a $n \times (p-1)$ design matrix, that excludes the ith node, and $\boldsymbol{\beta}$ is $(p-1) \times 1$ vector. The intercept is excluded, due to standardizing the data, so $\boldsymbol{\beta}^{(j)}$ contains $p-1$ regression coefficients. To be clear, $\boldsymbol{\beta}_j$ denotes the vector of coefficients for the jth regression model, where the individual elements are defined as β_{ij}. The residuals are assumed to follow $\boldsymbol{\varepsilon} \sim \mathcal{N}(0, \sigma_j^2)$,

where σ_j^2 is the residual variance for the jth node. The regression coefficients and error variances then correspond to the off-diagonal and diagonal elements of $\boldsymbol{\Theta}$–i.e.,

$$\theta_{ij} = \frac{-\beta_{ij}}{\sigma_j^2} \text{ and } \theta_{jj} = \frac{1}{\sigma_j^2}, \tag{16}$$

where θ_{ij} denotes the covariance corresponding to ith row and jth column of $\boldsymbol{\Theta}$. The diagonal of $\boldsymbol{\Theta}$ is then denoted θ_{jj}. Consequently, for all posterior samples, $s = 1, ..., S$,

$$\beta_{ij}^{(s)} = \frac{-\theta_{ij}^{(s)}}{\theta_{jj}^{(s)}} \text{ and } \sigma_j^{2(s)} = \frac{1}{\theta_{jj}^{(s)}} \tag{17}$$

results in the posterior distribution for each regression coefficient and residual variance. This correspondence allows for computing measures of predictive accuracy such as Bayesian R^2.

Bayesian R^2

In this section, I describe a procedure for computing R^2 for each node in the network. In a Bayesian setting, this cannot be computed with classical R^2 because the estimate will sometimes include impossible values (e.g., greater than 1). I use the method for computing Bayesian R^2 that was recently described in Gelman, Goodrich, Gabry, and Ali (2017). First define a variance function–i.e.,

$$V_{i=1}^N z_i = \frac{1}{N-1} \sum_{i=1}^N (z_i - \bar{z})^2, \tag{18}$$

where z is an arbitrary $N \times 1$ dimensional vector and N is the number of subjects in the sample. Thus i indexes the ith individuals score. This leads to the classical measure of variance explained being defined as

$$R^2 = \frac{V_{i=1}^N \hat{y}_n}{V_{i=1}^N y_n}, \tag{19}$$

where $\hat{\mathbf{y}}$ is a vector the predicted values that is of dimensions $N \times 1$ and $\mathbf{y}$ is the observed outcome of the same dimensions.

The extension to Bayesian models is mathematically equivalent to metrics commonly used in survival analysis (Choodari-Oskooei, Royston, & Parmar, 2012; Kent & O'Quigley, 1988). For each node or variable, instead of having a point estimate, there is a set of posterior draws for each regression coefficient that were obtained following Equation 16. Denote the posterior samples as $\boldsymbol{\theta}^{(s)}, s = 1, ..., S$. The predicted values are computed as $y_{j_i}^{(s)} = \mathrm{E}(\mathbf{y}_j | X_{(-j)}^*, \boldsymbol{\theta}^{*(s)})$, and the errors as $\epsilon_{j_i}^{(s)} = y_{j_i} - \hat{y}_{j_i}^s$. Here j refers to that specific node, i is the respective observation, and $*$ denotes the subset of selected predictors and posterior samples–e.g.,, $X_{(-j)}^* \subset X_{(-j)}$. This results in
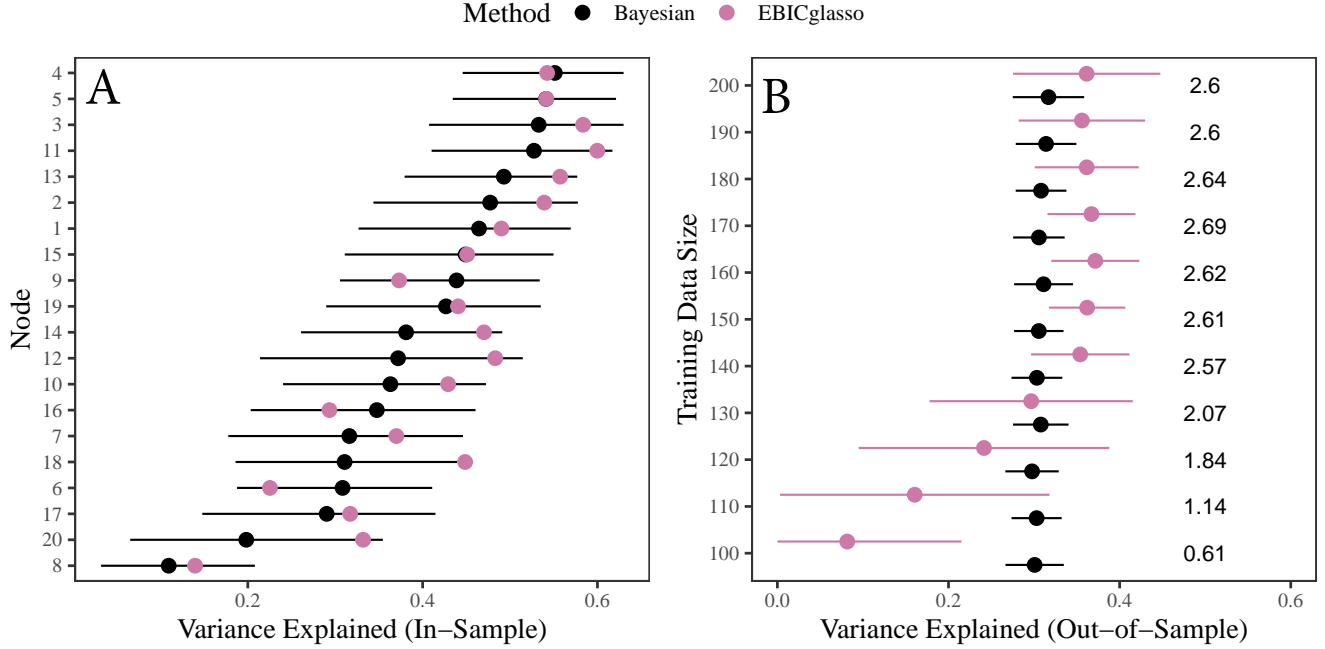
Figure 4. A) The panel includes in-sample variance explained for glasso$_{EBIC}$ and the Bayesian method. The networks for were selected, and then R^2 was computed for each node. The Bayesian methods readily provide a measure of uncertainty. The error bars correspond to 95 % credible intervals. The notable differences could be attributed to connectivity levels. The Bayesian networks were relatively sparse compared to those estimated with glasso$_{EBIC}$. I confirmed this by using a more liberal decision rule for the Bayesian method and the estimates converged to be very similar. B) This panel included out-of-sample variance explained. The numbers inside correspond to how many times more edges were selected by glasso$_{EBIC}$ than the Bayesian method. For example, 2.6 indicated that glasso$_{EBIC}$ had over twice as many connections in the network. Importantly, this plot shows that regularization can be advantageous for predicting new data (although the networks were relatively dense).

$$\text{Bayesian } R^2 = \frac{V_{i=1}^N \hat{y}_{j_i}^{(s)}}{V_{i=1}^N \hat{y}_{j_i}^{(s)} + V_{i=1}^N \epsilon_{j_i}^{(s)}}, \qquad (20)$$

which can be summarized in a variety of ways, for example the posterior mean, median or mode, as well as providing a measure of certainty. Note that, in contrast to Equation 19, Bayesian R^2 considers the predictive distribution in both the numerator and denominator. Further, this method can be extended to data that was not included in the original fitting. This provides a measure of *out-of-sample* variance explained.

Illustrative Example

In this section, I compute Bayesian R^2, both in- and out-of-sample, for each variable in the post-traumatic stress disorder data (Table 2). As a point of reference, I also include classical R^2 (Equation 19) that was computed from glasso$_{EBIC}$. This was accomplished by first selecting the network with the default settings of **qgraph**, and then each row of the selected precision matrix Θ was converted to the corresponding regression coefficients (Equation 16). For the

Bayesian method, the edge set was selected with posterior probabilities exceeding 0.975. This should correspond to an expected specificity of 0.95 (Table 1).

In-Sample Predictability. These results are reported in Figure 4 (panel A), where the estimates ordered from smallest to largest based on the Bayesian method. Additionally, for the Bayesian edges, I have included the distribution of R^2 estimates. There are again notable differences between glasso$_{EBIC}$ and BGGM, when relying exclusively on the point estimates, in that the order of importance (defined according to R^2) is not the same. That is, for BGGM node 4 (i.e., emotional cue reactivity) has the largest estimate, whereas the most important variable for glasso$_{EBIC}$ is apparently node 11 (i.e., negative trauma-related emotions). Of course, because the Bayesian method provides a measure of uncertainty, the differences in estimates can be evaluated. Here the difference in R^2 was 0.02 (95 % CrI = [-0.10, 0.15]), thereby indicating uncertainty with respect to the most important node in the network. This highlights an additional advantage of BGGM, in that accounting for uncertainty guards against over confident inferences. Further, not only was the ordering different, but it was also the case that one of the
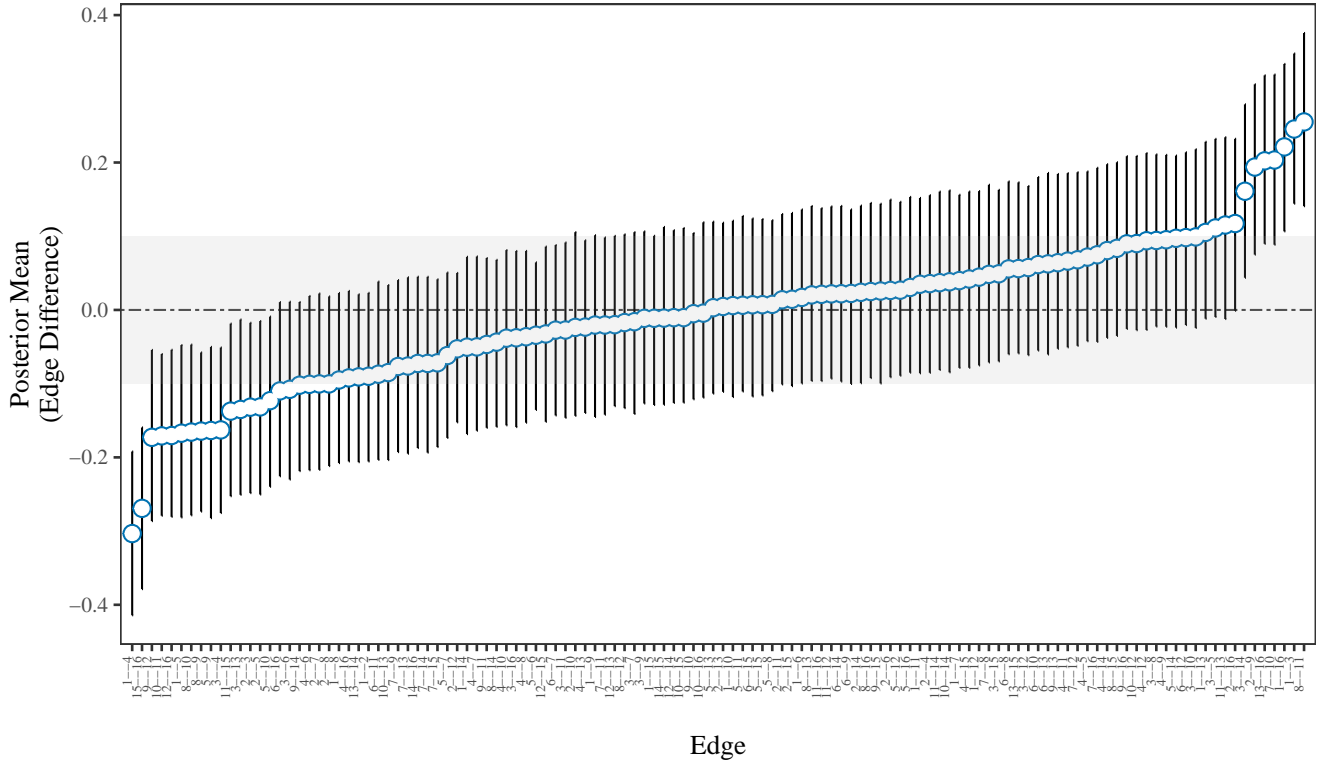
Figure 5. This figure includes edge differences between networks estimated from different groups (Table 3). The were computed by first estimating the edges for each group, and then the posterior distributions were subtracted. The grey area is the region of practical equivalence (±0.1). The error bars correspond to 99 % credible intervals. If an interval is *fully* contained in the rope, this would suggest that the edges were practically equivalent. Differences can be inferred either in reference to zero or the rope. For the former interval exclusion of zero would suggest a difference between networks. On the other hand, interval exclusion of the rope would suggest a practically meaningful difference. The rope must be decided based on subjective grounds, and perhaps is somewhat arbitrary, but I would argue no more than choosing a significance level.

glasso$_{EBIC}$ based estimates was not included in the Bayesian measure of variance explained (e.g., node 18). This could either suggest the Bayesian model is underfitting the data, based on the relatively sparse network (Figure 2), or that glasso$_{EBIC}$ is overfitting the data due to the dense graphical structure (Figure 2).

Out-of-Sample Predictability. In this section, I describe the out-of-sample R^2 results. I assumed varying sample sizes, in increments of 10, for the training data ($D_{\text{train}} = 100, ..., 200$). The sample sizes for the test data D_{test} was then $221 - D_{\text{train}}$. Importantly, this allowed for assessing predictability as a function of the p/n ratio. For example, with $D_{\text{train}} = 100$ this corresponds to $20/100 = 0.20$. For each training size, a subset of data was randomly selected, the network structure was selected, R^2 was computed on D_{test}, and then averaged across nodes. I was also interested in the number of variables selected which allowed for evaluating the trade-off between parsimony and predictive loss (Goutis, 1998; Piironen & Vehtari, 2017). R^2 and the sparsity levels were averaged across 100 simulation trials.

The out-of-sample R^2 results are provided in Figure 5 (panel B). The numbers inside the plot corresponds to how many times more edges were selected by glasso$_{EBIC}$ than the Bayesian method. For example, a value of 1 would indicate the graphs had the same number of connections. For the smaller training data sizes, the Bayesian method had better out-of-sample performance. This is possibly due to glasso$_{EBIC}$ selecting empty graphs. On the other hand, when the training data sizes increased, the regularized method regularly outperformed the nonregularized method. These difference were around 5 %. However, there is also a trade-off to consider. That is, if the goal is explicitly prediction glasso$_{EBIC}$ is advantageous, but a researcher must consider whether 5 % (approximately) more variance explained is worth nearly 2.5 times as many connections (i.e., a denser network)

Network Comparisons

The proposed methodology can also be extended to networks estimates from different groups. The customary ap-

proach is to use the R package **NetworkComparisonTest**. That approach employs ℓ_1-regularization and a permutation based significance test. One of the methods tests edge differences between networks. In a Bayesian setting, this can be accomplished by subtracting posterior distributions (Equation 11). Thus, for two networks, pairwise differences and equivalence can tested. The latter is not possible with **NetworkComparisonTest**. I applied the Bayesian method to post-traumatic stress disorder symptoms that were measured in two groups (N_{g1} = 926, and N_{g2} = 956). The symptoms and corresponding node numbers are provided in Table 3. Detailed information about the samples is provided in Fried et al. (2018).

Figure 5 includes the edge differences for the Bayesian methods. They are summarized with the posterior means and 99 % credible intervals. Further, I also shaded the area from ±0.1 which corresponds to the region of practical equivalence. Note that, while not shown here, **NetworkComparisonTest** does not provide a measure of uncertainty (e.g., a confidence interval) for the edge differences. Thus this highlights a key advantage of the Bayesian methods, in that a posterior distribution is provided for each edge difference. In this case, the directional posterior probabilities exceeded 99 % for several edge differences. This can be seen by noting the intervals that did not include zero. On the other hand, with a rope defined as a small effect size, none of the interval were fully included in the null region. Of course, it would be possible to change both the rope and the credible interval width.

Discussion

In this work, I introduced novel Bayesian methodology for Gaussian graphical models. The presented methods were specifically designed for low-dimensional settings ($p << n$), although they can also estimate regularized networks (Section: Regularized Estimation). I also provided methods that allow for extending inferences beyond the conditional (in)dependence structure. These methods are fully compatible with recently developed techniques in the Bayesian literature for computing R^2 and out-of-sample predictive accuracy, wherein the generalizability of the estimated network can be determined. This is particularly important, because there is an emerging interest in extending inference beyond explanation to also consider prediction in psychology (Yarkoni & Westfall, 2017). Further, although I specifically focused on the most common psychological applications, this work provides several novel contributions to the broader GGM literature which has relatively few Bayesian methods to date (Kuismin & Sillanpää, 2017).

With several illustrative examples, I contrasted Bayesian methodology and the current default approach in psychology (glasso$_{EBIC}$). Importantly, as previously mentioned (Introduction), ℓ_1-regularization only provides point estimates

which can present challenges for inference. For example, even when using a bootstrap strategy, the point mass at zero presents problems for computing valid confidence intervals (Bühlmann et al., 2014) and thus maintaining nominal error rates (Figure 3; panel B). The presented Bayesian method overcomes this issue–a full posterior distribution is provided for all quantities of interest such as the edge weights (Figure 5) and variance explained (Figure 4). These were obtained efficiently, as indicated by the timing comparison between methods (glasso$_{EBIC}$: 18.6 minutes vs. BGGM ≈ 0.5 seconds). This provides an opportunity for nuanced inferences in practical applications, in that accounting for uncertainty can guard against overconfidence that can arise when ranking variable importance according to the respective point estimates.

Importantly, the contrasts I made between methods was strictly for illustrative purposes. It is not possible to suggest on method is better than another from a sample of data (Jones, Williams, & McNally, 2019). Simulation would be needed for this purpose. As such, I emphasize that the advantageous presented here are not that the Bayesian methodology is necessarily providing a more accurate estimate, but rather there are key advantages in general. For example, the ability to compute the probability of null values.

For this work I did not draw sharp contrast from frequentist methods which has become commonplace when describing Bayesian methods in psychology (Kruschke, 2013). This is not a position I adopted for this work, where I used repeated sampling to validate the performance of the proposed decision rule (Figure: 1; Rubin, 1984). It should be noted that the commonly cited benefits of Bayesian methods are not readily apparent for GGMs, for example incorporating information via prior distributions (Zondervan-Zwijnenburg, Peeters, Depaoli, & Van de Schoot, 2017). In the context of GGMs, while it is possible to strategically use the prior distributions to provide shrinkage towards an identity matrix ($\rho_{ij} \to 0$; Kubokawa & Srivastava, 2008; Kuismin & Sillanpää, 2016; Leday & Richardson, 2018), incorporating question-specific prior information is not straightforward. This is due to the number of parameters being estimated, which makes specifying informative priors for each covariance a challenging task, as well as the positive-definite constraint on the covariance matrix (Joe, 2006). However, in my view, the greatest strength of Bayesian methods is the richness of information provided from drawing samples from the posterior distribution, in addition to what is being conditioned on. That is, conditional on the data and fitted model, Bayesian methods not only estimate probability distributions for the parameters, but can also provide a measure uncertainty for quantities of interest. It is this advantage, which I highlighted throughout this work, that provides a valuable contribution to the psychological network literature

Regularized Estimation

This method can readily be extended to regularized estimation. The Wishart distribution can provide a "ridge-type" estimate that can be achieved in closed form (Kubokawa & Srivastava, 2008; Kuismin & Sillanpää, 2016). For example, in reference to the MAP solution (Equation 6), regularization can be achieved by adjusting ϵ. However, it should be noted that the previously demonstrated calibration will be compromised (Table 1 and Figure 1). Not only is the point estimate shrunken towards zero, but the posterior can also become narrower which can then increase the false positive rate. Of course, if the explicit goal is prediction, then having too dense of a graph is not necessarily an issue and thus regularization would potentially offer some benefits in these situations. This can be seen in Figure 5, where glasso$_{EBIC}$ had excellent performance predicting unseen data.

Future Directions

There are limitations that can be addressed with future research. First, I did not perform extensive simulations to characterize the performance of the presented methodology. It should be noted that the dimensions of the most common psychological applications ($p < n$) do not present challenges for accurately determining the edge set (Williams & Rast, 2018). As such, considering that I demonstrated approximate calibration to specificity (Figure 1), examining less than ideal simulation conditions (e.g., violated assumptions) is an important future direction for both nonregularized and regularized method. Second, a commonly cited advantage of Bayesian methods is gaining evidence for the null hypothesis (Rouder, Speckman, Sun, Morey, & Iverson, 2009). However, the hypothesis being tested in this work provided evidence for the direction of the partial correlations or a null region. This cannot assess evidence in favor of the null hypothesis compared to some alternative hypothesis. This distinction is often referred to as Bayesian estimation vs. hypothesis testing. The methods in the work belong to the former, and thus, the prior distribution plays little role in inference. With colleagues, I have developed Bayesian methodology for hypothesis testing in GGMs (Williams & Mulder, 2019a; Williams, Rast, Pericchi, & Mulder, 2019). These methods are also implemented in the package BGGM.

Conclusion

Gaussian graphical models are becoming increasing popular tool in psychology. To date, the default estimation and inferential framework has been frequentist, relying upon methods developed for high-dimensional data. It is not only important to developed methods specifically for the most common psychological applications, but also approaches that utilize the advantages of Bayesian statistics. The presented Bayesian methodology was tailored for this purpose;

for example, by extending inference beyond the conditional (in)dependence structure. I have implemented these methods in the R package BGGM.

References

Alvarez, I., Niemi, J., & Simpson, M. (2014). Bayesian inference for a covariance matrix. In *Annual conference on applied statistics in agriculture* (Vol. 20, pp. 1669–1696). doi: 10.1214/aos/1176348885

Armour, C., Fried, E. I., Deserno, M. K., Tsai, J., & Pietrzak, R. H. (2017). A network analysis of DSM-5 posttraumatic stress disorder symptoms and correlates in U.S. military veterans. *Journal of Anxiety Disorders*, *45*(May 2013), 49–59. doi: 10.1016/j.janxdis.2016.11.008

Atay-Kayis, a., & Massam, H. (2005). A Monte Carlo method to compute the marginal likelihood in non decomposable graphical Gaussian models. *Biometrika*, *92*(2), 317–335.

Avagyan, V., Alonso, A. M., & Nogales, F. J. (2017, 10). Improving the Graphical Lasso Estimation for the Precision Matrix Through Roots of the Sample Covariance Matrix. *Journal of Computational and Graphical Statistics*, *26*(4), 865–872. doi: 10.1080/10618600.2017.1340890

Baba, K., Shibata, R., & Sibuya, M. (2004, 12). Partial Correlation and Conditional Correlation as Measures of Conditional Independence. *Australian & New Zealand Journal of Statistics*, *46*(4), 657–664. doi: 10.1111/j.1467-842X.2004.00360.x

Baba, K., & Sibuya, M. (2005). Equivalence of Partial and Conditional Correlation Coefficients. *J. Japan Statist. Soc.*, *35*(1), 1–19.

Barbieri, M. M., & Berger, J. O. (2004). Optimal predictive model selection. *Annals of Statistics*, *32*(3), 870–897. doi: 10.1214/009053604000000238

Barnard, J., McCulloch, R., & Meng, X.-L. (2000). Modelling Covariance Matrices in Terms of StandardDeviations and Correlations With Applications to Shrinkage. *Statistica Sinica*, *10*(4), 1281–1311. doi: 10.2307/24306780

Barrientos, A. F., & Peña, V. (2017). Bayesian Bootstraps for Massive Data.

Betancourt, M. (2018). Calibrating Model-Based Inferences and Decisions. *ArXiv preprint*, 1–35.

Birkeland, M. S., Blix, I., Solberg, , & Heir, T. (2017, 12). Gender differences in posttraumatic stress symptoms after a terrorist attack: A network approach. *Frontiers in Psychology*, *8*, 2091. doi: 10.3389/FPSYG.2017.02091

Borsboom, D., & Cramer, A. O. (2013). Network Analysis: An Integrative Approach to the Structure of Psychopathology. *Annual Review of Clinical Psychology*, *9*(1), 91–121. doi: 10.1146/annurev-clinpsy-050212-185608

Bühlmann, P., Kalisch, M., & Meier, L. (2014). High-Dimensional Statistics with a View Toward Applications in Biology. *Annual Review of Statistics and Its Application*, *1*(1), 255–278. doi: 10.1146/annurev-statistics-022513-115545

Burnham, K. P., & Anderson, D. R. (2004). Multimodel inference: Understanding AIC and BIC in model selection. *Sociological Methods and Research*, *33*(2), 261–304. doi: 10.1177/0049124104268644

Cai, T., Liu, W., & Luo, X. (2011). A constrained 1minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association*, *106*(494), 594–607. doi: 10.1198/jasa.2011.tm10155

Cao, X., Wang, L., Cao, C., Fang, R., Chen, C., Hall, B. J., & Elhai, J. D. (2018, 8). Sex differences in global and local connectivity of adolescent posttraumatic stress disorder symptoms. *Journal of child psychology and psychiatry, and allied disciplines*. doi: 10.1111/jcpp.12963

Casella, G., Girón, F. J., Martinez, M. L., & Moreno, E. (2009). Consistency of bayesian procedures for variable selection. *Annals of Statistics*, *37*(3), 1207–1228. doi: 10.1214/08-AOS606

Choodari-Oskooei, B., Royston, P., & Parmar, M. K. (2012, 10). A simulation study of predictive ability measures in a survival model I: Explained variation measures. *Statistics in Medicine*, *31*(23), 2627–2643. doi: 10.1002/sim.4242

Chopin, N., Gadat, S., Guedj, B., Guyader, A., & Vernet, E. (2015). ON SOME RECENT ADVANCES ON HIGH DIMENSIONAL BAYESIAN STATISTICS. , *51*, 293–319. doi: 10.1051/proc/201551016

Cohen, J., & Cohen, P. (1983). *Applied multiple regression/correlation analysis for the behavioral sciences* (2nd ed.). Hillsdale: L. Erlbaum Associates.

Dalege, J., Borsboom, D., van Harreveld, F., & van der Maas, H. L. J. (2017). A Network Perspective on Political Attitudes: Testing the Connectivity Hypothesis.

Das, A., Sampson, A., Lainscsek, C., Muller, L., Lin, W., Doyle, J., ... Sejnowski, T. (2017). Interpretation of the Precision Matrix and Its Application in Estimating Sparse Brain Connectivity during Sleep Spindles from Human Electrocorticography Recordings. *Neural computation*, *29*(3), 603–642. doi: 10.1162/NECO{_}a{_}00936

Dempster, A. (1972, 3). Covariance Selection. *Biometrics*, *28*(1), 157–175. doi: 10.2307/2528966

Deserno, M. K., Borsboom, D., Begeer, S., & Geurts, H. M. (2017). Multicausal systems ask for multicausal approaches: A network perspective on subjective well-being in individuals with autism spectrum disorder. *Autism*, *21*(8), 960–971. doi: 10.1177/1362361316660309

Dezeure, R., Bühlmann, P., Meier, L., & Meinshausen, N. (2014). High-Dimensional Inference: Confidence Intervals, p-Values and R-Software hdi. , 1–42. doi: 10.1214/15-STS527

Dobra, A., Lenkoski, A., & Rodriguez, A. (2010). Bayesian inference for general Gaussian graphical models with application to multivariate lattice data. , 1–30.

Epskamp, S. (2016). Regularized Gaussian Psychological Networks: Brief Report on the Performance of Extended BIC Model Selection. *arXiv*, 1–6.

Epskamp, S., Borsboom, D., & Fried, E. I. (2018). Estimating psychological networks and their accuracy: A tutorial paper. *Behavior Research Methods*, *50*(1), 195–212. doi: 10.3758/s13428-017-0862-1

Epskamp, S., & Fried, E. I. (2016). A Tutorial on Regularized Partial Correlation Networks. *arXiv*. doi: 10.1103/PhysRevB.69.161303

Epskamp, S., Kruis, J., & Marsman, M. (2017). Estimating psychopathological networks: Be careful what you wish for. *PLoS ONE*, *12*(6), 1–13. doi: 10.1371/journal.pone.0179891

Epskamp, S., Waldorp, L. J., Mottus, R., & Borsboom, D. (2018, 7). The Gaussian Graphical Model in Cross-Sectional and Time-Series Data. *Multivariate Behavioral Research*, *53*(4), 453–480. doi: 10.1080/00273171.2018.1454823

Fan, J., Liao, Y., & Liu, H. (2016). An overview of the estimation of large covariance and precision matrices. *Econometrics Journal*, *19*(1), C1-C32. doi: 10.1111/ectj.12061

Fried, E. I., Eidhof, M. B., Palic, S., Costantini, G., Dijk, H. M. H.-v., Bockting, C. L. H., ... Karstoft, K.-i. (2018). Replicability and Generalizability of Posttraumatic Stress Disorder (PTSD) Networks : A Cross-Cultural Multisite Study of PTSD Symptoms in Four Trauma Patient Samples. *Clinical Psychological Science*, *6*(3), 335–351. doi: 10.1177/2167702617745092

Friedman, J., Hastie, T., & Tibshirani, R. (2008, 7). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics*, *9*(3), 432–441. doi: 10.1093/biostatistics/kxm045

Gelman, A., Carlin, J. B., Stern, H. S., Dunson, B. D., Vehtari, A., & Rubin, D. B. (2014). *Bayesian Data Analysis*. Boca Raton: CRC Press.

Gelman, A., Goodrich, B., Gabry, J., & Ali, I. (2017). R-squared for Bayesian regression models.

George, E. I., & McCulloch, R. E. (1993). Variable selection via Gibbs sampling. *Journal of the American Statistical Association*, *88*(August 2015), 881–889. doi: 10.1080/01621459.1993.10476353

Goutis, C. (1998). Model choice in generalised linear models: a Bayesian approach via Kullback-Leibler projections. *Biometrika*, *85*(1), 29–37. doi: 10.1093/biomet/85.1.29

Green, P. J. (1995). Reversible Jump Markov Chain Monte Carlo Computation and Bayesian Model Determination. *Biometrika*, *82*(4), 711–732. doi: 10.2307/2337340

Gutiérrez-Peña, E., Smith, A. F. M., Bernardo, J. M., Consonni, G., Veronese, P., George, E. I., ... Morris, C. N. (1997, 6). Exponential and bayesian conjugate families: Review and extensions. *Test*, *6*(1), 1–90. doi: 10.1007/BF02564426

Hajdúk, M., Klein, H. S., Harvey, P. D., Penn, D. L., & Pinkham, A. E. (2018, 7). Paranoia and interpersonal functioning across the continuum from healthy to pathological - Network analysis. *British Journal of Clinical Psychology*. doi: 10.1111/bjc.12199

Haslbeck, J. M. B., & Fried, E. I. (2017). How predictable are symptoms in psychopathological networks? A reanalysis of 18 published datasets. *Psychological Medicine*, *47*(16), 2267–2276. doi: 10.1017/S0033291717001258

Haslbeck, J. M. B., & Waldorp, L. J. (2018, 4). How well do network models predict observations? On the importance of predictability in network models. *Behavior Research Methods*, *50*(2), 853–861. doi: 10.3758/s13428-017-0910-x

Hastie, T., Tibshirani, R., & Wainwright, M. (2015). *Statistical Learning with Sparsity: The Lasso and Generalizations*. Boca Raton: CRC Press. doi: 10.1201/b18401-1

Hoeting, J. A., Madigan, D., Raftery, A. E., & Volinsky, C. T. (1999). Bayesian Model Averaging: A Tutorial. *Statistical Science*, *14*(4), 382–417. doi: 10.2307/2676803

Højsgaard, S., Edwards, D., & Lauritzen, S. (2012). *Graphical Models with R*. doi: 10.1007/978-1-4614-2299-0

Isvoranu, A. M., Van Borkulo, C. D., Boyette, L. L., Wigman, J. T., Vinkers, C. H., Borsboom, D., . . . Myin-Germeys, I. (2017). A network approach to psychosis: Pathways between childhood trauma and psychotic symptoms. *Schizophrenia Bulletin*, *43*(1), 187–196. doi: 10.1093/schbul/sbw055

Janková, J., & van de Geer, S. (2015). Confidence intervals for high-dimensional inverse covariance estimation. *Electronic Journal of Statistics*, *9*, 1205–1229. doi: 10.1214/15-EJS1031

Janková, J., & van de Geer, S. (2017). Honest confidence regions and optimality in high-dimensional precision matrix estimation. *Test*, *26*(1), 143–162. doi: 10.1007/s11749-016-0503-5

Joe, H. (2006). Generating random correlation matrices based on partial correlation vines and the onion method. *Journal of Multivariate Analysis*, *97*(10), 2177–2189. doi: 10.1016/j.jmva.2005.05.010

Jones, P. J., Williams, D. R., & McNally, R. J. (2019). Sampling Variability is not Nonreplication: A Bayesian Reanalysis of Forbes, Wright, Markon, & Krueger. *PsyArXiv*. doi: 10.31234/OSF.IO/EGWFJ

Kent, J. T., & O'Quigley, J. (1988, 9). Measures of dependence for censored survival data. *Biometrika*, *75*(3), 525–534. doi: 10.1093/biomet/75.3.525

Khondker, Z. S., Zhu, H., Chu, H., Lin, W., & Ibrahim, J. G. (2013, 4). The Bayesian Covariance Lasso. *Statistics and its interface*, *6*(2), 243–259. doi: 10.1007/s11103-011-9767-z.Plastid

Kruschke, J. K. (2011). Bayesian Assessment of Null Values Via Parameter Estimation and Model Comparison. *Perspectives on Psychological Science*, *6*, 299–312. doi: 10.1177/1745691611406925

Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, *142*(2), 573–603. doi: 10.1037/a0029146

Kubokawa, T., & Srivastava, M. S. (2008). Estimation of the precision matrix of a singular Wishart distribution and its application in high-dimensional data. *Journal of Multivariate Analysis*, *99*(9), 1906–1928. doi: 10.1016/j.jmva.2008.01.016

Kuismin, M., & Sillanpää, M. (2017). Estimation of covariance and precision matrix, network structure, and a view toward systems biology. *Wiley Interdisciplinary Reviews: Computational Statistics*, *9*(6), 1–13. doi: 10.1002/wics.1415

Kuismin, M., & Sillanpää, M. J. (2016). Use of Wishart prior and simple extensions for sparse precision matrix estimation. *PLoS ONE*, *11*(2), e0148171. doi: 10.1371/journal.pone.0148171

Kwan, C. C. Y. (2014). A Regression-Based Interpretation of the Inverse of the Sample Covariance Matrix. *Spreadsheets in Education*, *7*(1).

Lauritzen, S. L. (1996). *Graphical models* (Vol. 17). Clarendon Press.

Leday, G. G. R., & Richardson, S. (2018). Fast Bayesian inference in large Gaussian graphical models. , 1–26.

Lenkoski, A. (2013). A direct sampler for G-Wishart variates. *Stat*, *2*(1), 119–128. doi: 10.1002/sta4.23

Liu, H., Zhang, Z., & Grimm, K. J. (2016, 5). Comparison of Inverse Wishart and Separation-Strategy Priors for Bayesian Estimation of Covariance Parameter Matrix in Growth Curve Analysis. *Structural Equation Modeling*, *23*(3), 354–367. doi: 10.1080/10705511.2015.1057285

Lockhart, R., Taylor, J., Tibshirani, R. J., & Tibshirani, R. (2014). A significance test for the lasso. *Annals of Statistics*, *42*(2), 413–468. doi: 10.1214/13-AOS1175

Luo, X. (2014). A Hierarchical Graphical Model for Big Inverse Covariance Estimation with an Application to fMRI. , 1–21.

Mao, L., Van Hemert, J. L., Dash, S., & Dickerson, J. A. (2009). Arabidopsis gene co-expression network and its functional modules. *BMC Bioinformatics*, *10*(1), 346. doi: 10.1186/1471-2105-10-346

Mazumder, R., & Hastie, T. (2012, 3). Exact Covariance Thresholding into Connected Components for Large-Scale Graphical Lasso. *Journal of machine learning research : JMLR*, *13*, 781–794.

McNally, R. J., Heeren, A., & Robinaugh, D. J. (2017, 11). A Bayesian network analysis of posttraumatic stress disorder symptoms in adults reporting childhood sexual abuse. *European Journal of Psychotraumatology*, *8*(sup3), 1341276. doi: 10.1080/20008198.2017.1341276

McNally, R. J., Robinaugh, D. J., Wu, G. W. Y., Wang, L., Deserno, M. K., & Borsboom, D. (2015, 11). Mental Disorders as Causal Systems. *Clinical Psychological Science*, *3*(6), 836–849. doi: 10.1177/2167702614553230

McNeish, D. M. (2015). Using Lasso for Predictor Selection and to Assuage Overfitting: A Method Long Overlooked in Behavioral Sciences. *Multivariate Behavioral Research*, *50*(5), 471–484. doi: 10.1080/00273171.2015.1036965

Meinshausen, N., & Bühlmann, P. (2006). High-dimensional graphs and variable selection with the Lasso. *Annals of Statistics*, *34*(3), 1436–1462. doi: 10.1214/009053606000000281

Mohammadi, A., & Wit, E. C. (2015). Bayesian structure learning in sparse Gaussian graphical models. *Bayesian Analysis*, *10*(1), 109–138. doi: 10.1214/14-BA889

Newton, M., Polson, N. G., & Xu, J. (2018, 3). Weighted Bayesian Bootstrap for Scalable Bayes. Retrieved from http://arxiv.org/abs/1803.04559

Neyman, J. (1935, 9). On the Problem of Confidence Intervals. *The Annals of Mathematical Statistics*, *6*(3), 111–116. doi: 10.1214/aoms/1177732585

Nuijten, M. B., Wetzels, R., Matzke, D., Dolan, C. V., & Wagenmakers, E.-J. (2015, 3). A default Bayesian hypothesis test for mediation. *Behavior Research Methods*, *47*(1), 85–97. doi: 10.3758/s13428-014-0470-2

O'Hara, R. B., & Sillanpää, M. J. (2009). A review of bayesian variable selection methods: What, how and which. *Bayesian Analysis*, *4*(1), 85–118. doi: 10.1214/09-BA403

Pearl, J. (2009). Causal inference in statistics: An overview. *Statistics Surveys*, *3*, 96–146. doi: 10.1214/09-SS057

Piironen, J., & Vehtari, A. (2017). Comparison of Bayesian predictive methods for model selection. *Statistics and Computing*, *27*(3), 711–735. doi: 10.1007/s11222-016-9649-y

Rhemtulla, M., Fried, E. I., Aggen, S. H., Tuerlinckx, F., Kendler, K. S., & Borsboom, D. (2016). Network analysis of substance abuse and dependence symptoms. *Drug and Alcohol Dependence*, *161*, 230–237. doi: 10.1016/j.drugalcdep.2016.02.005

Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: parameter estimation and Bayes factors. *Psychonomic Bulletin and Review*, *25*(1), 102–113. doi: 10.3758/s13423-017-1420-7

Rouder, J. N., Speckman, P. L., Sun, D., Morey, R. D., & Iverson, G. (2009). Bayesian t tests for accepting and rejecting the null hypothesis. *Psychonomic Bulletin and Review*, *16*(2), 225–237. doi: 10.3758/PBR.16.2.225

Rubin, D. B. (1984). Bayesianly Justifiable and Relevant Frequency Calculations for the Applied Statistician. *The Annals of Statistics*, *12*(4), 1151–1172. doi: 10.1214/aos/1176346785

Schuurman, N. K., Grasman, R. P., & Hamaker, E. L. (2016, 5). A Comparison of Inverse-Wishart Prior Specifications for Covariance Matrices in Multilevel Autoregressive Models. *Multivariate Behavioral Research*, *51*(2-3), 185–206. doi: 10.1080/00273171.2015.1065398

Stephens, G. (1998). On the Inverse of the Covariance Matrix in Portfolio Analysis. *The Journal of Finance*, *53*(5), 1821–1827.

Taylor, J., & Tibshirani, R. (2017). Post-selection inference for 1-penalized likelihood models. *Canadian Journal of Statistics*(2015), 1–26. doi: 10.1002/cjs.11313

Tibshirani, R. (1996). *Regression Selection and Shrinkage via the Lasso* (Vol. 58) (No. 1). doi: 10.2307/2346178

Tibshirani, R. J., Taylor, J., Lockhart, R., & Tibshirani, R. (2016). Exact Post-Selection Inference for Sequential Regression Procedures. *Journal of the American Statistical Association*, *111*(514), 600–620. doi: 10.1080/01621459.2015.1108848

Van De Geer, S., Bühlmann, P., Ritov, Y., & Dezeure, R. (2014). On asymptotically optimal confidence regions and tests for high-dimensional models. *Annals of Statistics*, *42*(3), 1166–1202. doi: 10.1214/14-AOS1221

van Rooijen, G., Isvoranu, A. M., Meijer, C. J., van Borkulo, C. D., Ruhé, H. G., & de Haan, L. (2017). A symptom network structure of the psychosis spectrum. *Schizophrenia Research*, *189*, 75–83. doi: 10.1016/j.schres.2017.02.018

Wagenmakers, E. J., Love, J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., … Morey, R. D. (2018). Bayesian inference for psychology. Part II: Example applications with JASP. *Psychonomic Bulletin and Review*, *25*(1), 58–76. doi: 10.3758/s13423-017-1323-7

Wagenmakers, E. J., Marsman, M., Jamil, T., Ly, A., Verhagen, J., Love, J., … Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages and practical ramifications. *Psychonomic Bulletin and Review*, *25*(1), 35–57. doi: 10.3758/s13423-017-1343-3

Wang, H. (2012). Bayesian graphical lasso models and efficient posterior computation. *Bayesian Analysis*, *7*(4), 867–886. doi: 10.1214/12-BA729

Wetzels, R., & Wagenmakers, E.-J. (2012, 12). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic bulletin & review*, *19*(6), 1057–64. doi: 10.3758/s13423-012-0295-x

Whittaker, J. (1990). *Graphical models in applied multivariate statistics*. New York: Wiley.

Wille, A., Zimmermann, P., Vranova, E., Furholz, A., Laule, O., Bleuler, S., … Bühlmann, P. (2004). Sparse graphical Gaussian modeling of the isoprenoid gene network in Arabidopsis thaliana. *Genome Biology*, *5*(11), -. doi: 10.1186/gb-2004-5-11-r92

Williams, D. R., & Mulder, J. (2019a). Bayesian Hypothesis Testing for Gaussian Graphical Models: Conditional Independence and Order Constraints.

Williams, D. R., & Mulder, J. (2019b). BGGM : A R Package for Bayesian Gaussian Graphical Models.

Williams, D. R., & Rast, P. (2018). Back to the basics: Rethinking partial correlation network methodology. , 1–15. doi: 10.17605/OSF.IO/FNDRU

Williams, D. R., Rast, P., Pericchi, L. R., & Mulder, J. (2019). Comparing Gaussian Graphical Models with the Posterior Predictive Distribution and Bayesian Model Selection. doi: https://doi.org/10.31234/osf.io/yt386

Williams, D. R., Rhemtulla, M., Wysocki, A. C., & Rast, P. (2018). On Non-Regularized Estimation of Psychological Networks.

Williams, D. R., Rhemtulla, M., Wysocki, A. C., & Rast, P. (2019). On Nonregularized Estimation of Psychological Networks. *Multivariate Behavioral Research*, *0*(0), 1–23. doi: 10.1080/00273171.2019.1575716

Yarkoni, T., & Westfall, J. (2017, 11). Choosing Prediction Over Explanation in Psychology: Lessons From Machine Learning. *Perspectives on Psychological Science*, *12*(6), 1100–1122. doi: 10.1177/1745691617693393

Zhang, C.-H., & Zhang, S. S. (2014). Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, *76*(1), 217–242. doi: 10.1111/rssb.12026

Zhao, P., & Yu, B. (2006). On Model Selection Consistency of Lasso. *The Journal of Machine Learning Research*, *7*, 2541–2563. doi: 10.1109/TIT.2006.883611

Zondervan-Zwijnenburg, M., Peeters, M., Depaoli, S., & Van de Schoot, R. (2017, 10). Where Do Priors Come From? Applying Guidelines to Construct Informative Priors in Small Sample Research. *Research in Human Development*, *14*(4), 305–320. doi: 10.1080/15427609.2017.1370966

Table 2
Node descriptions

Node	Symptom
1	Intrusive Thoughts
2	Nightmares
3	Flashbacks
4	Emotional cue reactivity
5	Psychological cue reactivity
6	Avoidance of thoughts
7	Avoidance of reminders
8	Trauma-related amnesia
9	Negative beliefs
10	Blame of self or others
11	Negative trauma-related emotions
12	Loss of interest
13	Detachment
14	Restricted affec
15	Irritability/anger
16	Self-destructive/reckless behavior
17	Hypervigilance
18	Exaggerated startle response
19	Difficulty concentrating
20	Sleep disturbance

Table 3
Node descriptions

Node	Symptom
1	Intrusive Thoughts
2	Nightmares
3	Flashbacks
4	Physiological/psychological reactivity
5	Avoidance of thoughts
6	Avoidance of situations
7	Amnesia
8	Disinterest in activities
9	Feeling detached
10	Emotional numbing
11	Foreshortened future
12	Sleep problems
13	Irritability
14	Concentration problems
15	Hypervigilance
16	Startle response