Bayesian Hypothesis Testing for Gaussian Graphical Models: Conditional Independence and Order Constraints

Donald R. Williams[a,*], Joris Mulder[bc]

[a]*Department of Psychology, University of California, Davis, USA*
[b]*Department of Methodology and Statistics, Tilburg University, The Netherlands*
[c]*Jheronimus Academy of Data Science, 's Hertogenbosch, The Netherlands*

Abstract

Gaussian graphical models (GGM; partial correlation networks) have become increasingly popular in the social and behavioral sciences for studying conditional (in)dependencies between variables. In this work, we introduce exploratory and confirmatory Bayesian tests for partial correlations. For the former, we first extend the customary GGM formulation that focuses on conditional dependence to also consider the null hypothesis of conditional independence for each partial correlation. Here a novel testing strategy is introduced that can provide evidence for a null, negative, or positive effect. We then introduce a test for hypotheses with order constraints on partial correlations. This allows for testing theoretical and clinical expectations in GGMs. The novel matrix$-F$ prior distribution is described that provides increased flexibility in specification compared to the Wishart prior. The methods are applied to PTSD symptoms. In several applications, we demonstrate how the exploratory and confirmatory approaches can work in tandem: hypotheses are formulated from an initial analysis and then tested in an independent dataset. The methodology is

---

*Corresponding Author
Email address: `drwwilliams@ucdavis.edu` (Donald R. Williams)

implemented in the R package **BGGM** .

Keywords: Gaussian graphical model, Bayes factor, partial correlation, order-constraints, matrix$-F$ prior

---

1. Introduction

Gaussian graphical models (GGM, aka partial correlation networks) have become increasingly popular in psychological science (Epskamp et al., 2018). The underlying rationale is that a group of variables, say, items from a psychometric scale, are thought to measure a dynamic system that mutually influence and interact with one another (Borsboom and Cramer, 2013). This is distinct from customary approaches, such as structural equation models, where observed items are indicators that are thought to share a common cause represented by an unobserved or latent variable. The items are usually assumed to be independent, conditionally on the latent variable. This perspective has been criticized by network proponents because it ignores the possibility of functional associations among observed variables (Borsboom and Cramer, 2013; Cramer et al., 2010). Exploring these relations has shed new light upon a variety of research areas including personality (e.g., Costantini et al., 2015) and psychopathology (e.g., Di Pierro et al., 2019; Werner et al., 2018).

In GGMs, the primary objective is to characterize conditional dependencies between pairs of variables. This is typically accomplished by determining which off-diagonal elements in the inverse of the covariance matrix (i.e., the precision matrix) are non-zero. This is referred to as *covariance selection* (Dempster, 1972; Peng et al., 2009). When these elements are standardized and the sign reversed, this results in partial correlations (ρ) that are pairwise relations controlling for all other variables in the model (Baba and Sibuya,

2005; Baba et al., 2004). The non-zero partial correlations constitute the underlying network structure of conditionally dependent effects. In a graph, they are represented visually as "connections" that link variables (Jones et al., 2018).

This work fills two gaps in the network literature. First, our methodology allows researchers to formally test the null hypothesis of conditional *independence* for each relation. Partial correlation networks are thought to represent causal skeletons (but see Ryan et al., 2019; Dablander and Hinne, 2019) and an important aspect of causal inference is conditional independence (Pearl, 2009). This requires assessing evidence for the null hypothesis. However, the typical approach merely detects non-zero effects and those not detected are set to zero (see for example, Drton and Perlman, 2004; Williams and Rast, 2019). Secondly, we introduce a method for confirmatory testing. GGMs are often estimated with data-driven procedures that do not readily allow for incorporating *a priori* expectations (e.g., see all methods in Table 1 of Kuismin and Sillanpää, 2017). In network theory, however, it has been argued that psychopathology symptoms "do not function as protective factors in the development of other symptoms" (p. 6, Borsboom et al., 2011). Hence, all of the partial correlations in a given network are often expected to be positive. This one-sided or order-restriction corresponds to the theoretical idea of a positive manifold (Horn and Cattell, 1966). Note here that testing such expectations directly implies more statistical power compared to less precise hypotheses, also from a Bayesian perspective (e.g., Mulder and Raftery, 2019). Additionally, there is now a wealth of networks analyses in the literature and synthesizing this information into formal theories to then test is a pressing challenge (Haslbeck et al., 2019). Together, these innovations are particularly useful for psychological applications.

Our methodology is built around exploratory and confirmatory Bayesian testing in GGMs. The former is specifically geared towards individual partial correlations, including two-sided and one-sided testing, as well as an exhaustive approach (i.e., $\mathcal{H}_0 : \rho = 0$ vs. $\mathcal{H}_1 : \rho > 0$ vs. $\mathcal{H}_2 : \rho < 0$). The confirmatory aspect builds upon Mulder (2016) and allows for testing competing sets of order constraints on multiple partial correlations. Network theory may suggest that a set of partial correlations, say, $\boldsymbol{\rho}_A$, is expected to be larger than another set of partial correlations, say, $\boldsymbol{\rho}_B$, which are all expected to be positive. This can be translated to an order hypothesis $\mathcal{H}_1 : \boldsymbol{\rho}_a > \boldsymbol{\rho}_b > \mathbf{0}$ that can be tested against competing order hypotheses or its complement, "not $\mathcal{H}_1$," for instance. We are not aware of any work in the classical literature for testing order hypotheses on partial correlations, although tests have been proposed for equality constrained null hypothesis testing of bivariate correlations (e.g., Krishnamoorthy and Xia, 2007; Preacher, 2006; Steiger, 1980). Furthermore, while our extension to partial correlations is novel, Bayes factors have been developed for testing order-constrained hypotheses in ANOVA (Klugkist et al., 2005), regression coefficients (Braeken et al., 2015), group variances (Böing-Messing et al., 2017; Dablander et al., 2020), bivariate correlations (Mulder, 2016), intraclass correlations (Mulder and Fox, 2013, 2018), network autocorrelation models (Dittrich et al., 2019, 2020), and multinomial models (Heck and Davis-Stober, 2019). The present work adds to this growing body of literature.

There are a variety of Bayesian approaches in the GGM literature, although they remain uncommon in practice (reviewed in Kuismin and Sillanpää, 2017). In the psychological literature, a Bayesian hypothesis test for one partial correlation has been developed under a linear regression model (Wetzels and Wagenmakers, 2012; Mulder et al., 2012).

In the broader GGM literature, there are several methods that primarily focus on conditionally dependent relations ($\rho \neq 0$), including with the Bayesian information criterion (Kuismin and Sillanpää, 2016), posterior inclusion probabilities (Mohammadi and Wit, 2015; Bhadra and Mallick, 2013), credible interval exclusion of zero (Khondker et al., 2013; Li et al., 2017), predictive utility (Williams et al., 2018), and Bayes factors (Giudici, 1995). Perhaps the most notable of these approaches employs the Wishart distribution (Kuismin and Sillanpää, 2016; Leday and Richardson, 2018; Tsukuma, 2014), or a generalization thereof (e.g., G-Wishart, Mohammadi and Wit, 2015), which is the conjugate prior distribution for the precision matrix (Gutiérrez-Peña et al., 1997). These methods have limitations of their own, for example, none allow for testing one-sided constraints. Furthermore, the customary Wishart distribution has been criticized for being too restrictive (e.g. the entire precision matrix is governed by only two parameters, Leonard and Hsu, 1992; Hsu et al., 2012). In this work, we overcome this issue by testing partial correlations with the recently introduced matrix$-F$ prior.

This paper is organized as follows. First notation specific to GGMs is introduced in Section 2. In Section 3, we introduce the method for exploratory testing for each partial correlation in the network. Here the focus is conditional independence, exhaustive testing, and prior specification. The confirmatory aspect of this work is then described in Section 4. Here the importance of testing order-constrained hypotheses in psychological networks is emphasized. The paper ends with a short discussion in Section 5.

2. The Gaussian Graphical Model

For multivariate normal data (Baba et al., 2004; Baba and Sibuya, 2005), a GGM captures conditional relations that are typically visualized to infer the underlying dependence structure (i.e., the partial correlation "network"; Højsgaard et al., 2012; Lauritzen, 1996). There is an undirected graph that is denoted $G = (V, E)$, consists of a vertex $V = \{1, \ldots, p\}$ and edge set $E \subset V \times V$. The former refers to "nodes" that are, say, symptoms in a psychopathology scale, whereas E is the estimated network structure. Let $\mathbf{y} = (y_1, \ldots, y_p)'$ be a random vector indexed by the graph's vertices that is assumed to follow a multivariate normal distribution, $\mathbf{y} \sim \mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with the mean vector $\boldsymbol{\mu} = (0_1, \ldots, 0_p)'$ and a $p \times p$ positive definite covariance matrix $\boldsymbol{\Sigma}$. In the following, we use $\mathbf{Y}$ to denote the $n \times p$ data matrix, where each row corresponds to the observations from the ith individual $i = \{1, \ldots, n\}$.

The undirected graph is obtained by determining which off-diagonal elements in the precision matrix, $\boldsymbol{\Theta} = \boldsymbol{\Sigma}^{-1}$, are non-zero. That is, $(i, j) \in E$ when node i and j are determined to be conditionally dependent and set to zero otherwise. One option is to compute credible intervals for θ_{ij}. This has an analytic solution when using the Wishart prior for $\boldsymbol{\Theta}$ (Equation 7 in Williams, 2018). Alternatively, standardizing $\boldsymbol{\Theta}$ and reversing the sign yields partial correlations, that is,

$$\rho_{ij \cdot z} = \frac{-\theta_{ij}}{\sqrt{\theta_{ii}\theta_{jj}}}, \tag{1}$$

where z contains the nodes conditioned on. To simplify the notation, we now denote the partial correlations with ρ_{ij}. Hence, it is possible to determine E by testing each partial

correlation in isolation (i.e., $\frac{1}{2}p(p-1)$ hypothesis tests). We adopt this strategy, which is then used to construct the $p \times p$ adjacency matrix, $\mathbf{A}^{CD}$, that is defined as

$$
\mathbf{A}_{ij}^{CD} = \begin{cases} 1 & \text{if } \rho_{ij} \neq 0 \\ 0 & \text{otherwise,} \end{cases}
\tag{2}
$$

where $\text{diag}(\mathbf{A}^{CD}) = 0$. The position of 1's in a given row corresponds to the "neighborhood" for node j (Meinshausen and Bühlmann, 2006) and they encode the conditional dependence structure (i.e., CD in 2). Although this is the customary formulation for GGMs, it is important to note that evidence for conditional independence (i.e., $\rho_{ij} = 0$) is not considered. This has resulted in some confusion in the psychological literature, in that nodes *not* sharing an edge are sometimes misinterpreted as conditionally independent. Our innovation is to define an additional $p \times p$ adjacency matrix, $\mathbf{A}^{CI}$, that encodes the conditional independence structure. This is similarly defined as

$$
\mathbf{A}_{ij}^{CI} = \begin{cases} 1 & \text{if } \rho_{ij} = 0 \\ 0 & \text{otherwise.} \end{cases}
\tag{3}
$$

Both (2) and (3) are determined with our novel Bayesian test introduced below. Hence, both adjacency matrices include relations that were supported by the data. Note that (2) can also include edges corresponding to an order-constrained hypothesis, say, $\mathcal{H}_1 : \rho_{ij} > 0$, which results in a structure with only positive relations. Furthermore, by characterizing both the conditional dependence and independence structures, this leads to an "ambigu-

7

ous" network (there was not evidence for or against an edge). That is, $\mathbf{1}_{pxp} - \mathbf{I}_p - \mathbf{A}^{CI} - \mathbf{A}^{CD} = \mathbf{A}^A$, where $\mathbf{1}_{pxp}$ is a $p \times p$ matrix of 1's and thus the 1's in $\mathbf{A}^A$ denote the ambiguous relations. The end result is that each partial correlations is classified into one of the three adjacency matrices.

To summarize, we have extended the customary GGM formulation to reflect the theoretical relevance of conditional independence. As an added bonus, there is a third adjacency matrix, $\mathbf{A}^A$, that guards against over confident inferences by indicating the nodes for which there was inconclusive evidence. These extensions beyond $\mathbf{A}^{CD}$ are major contributions.[1]

3. Exploratory Hypothesis Testing

3.1. Conditional independence and dependence

We determine the conditional independence and dependence structures (Equations 2 and 3) by comparing an equality constrained (null) hypothesis versus an unconstrained hypothesis

$$\mathcal{H}_0 : \rho_{ij} = 0 \tag{4}$$

$$\mathcal{H}_u : \rho_{ij} \in (-1, 1).$$

This resembles a classical testing framework, although, in this case, we are not assuming the null hypothesis of conditional independence is true. Rather, relative evidence for the

---

[1]Note that assessing null values is also possible with classical statistics. However, we are not aware of any examples doing so in Gaussian graphical models.

competing hypotheses is quantified for each partial correlation. One null hypothesis testing strategy employs a likelihood ratio test (e.g., pp 87 - 88, Højsgaard et al., 2012). This is informative because it indicates that we are testing nested models in (4). As a result, we use the Savage-Dickey ratio which simplifies computing the Bayes factor (Dickey, 1971; Wagenmakers et al., 2010). The hypothesis test in favor of the null hypothesis can be formulated as

$$
\begin{aligned}
\text{BF}_{0u} &= \frac{p(\mathbf{Y}|\mathcal{H}_0)}{p(\mathbf{Y}|\mathcal{H}_u)} \\
&= \frac{\int f_0(\mathbf{Y}|\Sigma_0, \mathcal{H}_0)\pi_0(\Sigma_0|\mathcal{H}_0)d\Sigma_0}{\int f_u(\mathbf{Y}|\Sigma_u, \mathcal{H}_u)\pi_u(\Sigma_u|\mathcal{H}_u)d\Sigma_u},
\end{aligned}
\tag{5}
$$

where $f_i(\mathbf{Y}|\Sigma_i, \mathcal{H}_i)$ and $\pi_i(\Sigma_i|\mathcal{H}_i)$ are the likelihood functions and the prior of the parameters under each hypothesis $i \in \{0, 1\}$. Furthermore, Σ_0 denotes a covariance matrix with $\theta_{ij} = 0$. This clarifies the relationship to the likelihood ratio test, as the full model is compared to a nested model with an off-diagonal element of $\mathbf{\Theta}$ constrained to zero (p. 87, Højsgaard et al., 2012). This general formulation was used in Leday and Richardson (see Equation 4, 2018), which extended the approach described in Giudici (1995) to high-dimensional settings ($n < p$). In Leday and Richardson (2018), an analytic expression was introduced for two-sided testing. However, the resulting Bayes factor was not scale-invariant and required rescaling. This is a result of testing $\mathcal{H}_0 : \theta_{ij} = 0$ vs. $\mathcal{H}_u : \theta_{ij} \neq 0$ (p. 5, Leday and Richardson, 2018), as opposed to the scale-invariant partial correlation in (4). Furthermore, focusing on $\mathbf{\Theta}$ presents challenges for prior specification. This is also due to lacking scale-invariance.

Our analytic and sampling-based solutions target the partial correlations in isolation of one another, that is,

$$\text{BF}_{0u} = \frac{p(\mathbf{Y}|\mathcal{H}_0)}{p(\mathbf{Y}|\mathcal{H}_u)} = \frac{p(\rho_{ij} = 0|\mathbf{Y}, \mathcal{H}_u)}{p(\rho_{ij} = 0|\mathcal{H}_u)}, \tag{6}$$

where the off-diagonal elements in $\mathbf{\Theta}$ have been standardized (Equation 1). In words, by only considering $\mathcal{H}_u$ with respect to ρ_{ij}, the Bayes factor can be computed as the unconstrained posterior density evaluated at zero divided by the prior density also evaluated at zero. This is the Savage-Dickey ratio (Dickey, 1971; Wagenmakers et al., 2010; Mulder et al., 2012). This expression of the Bayes factor holds under certain prior conditions (Dickey, 1971; Heck, 2018), which will be discussed later. In the psychological literature, an analogous approach has been used for both correlations (Wagenmakers et al., 2016; Marsman and Wagenmakers, 2017) and one partial correlation (Nuijten et al., 2015). The advantage of (6) is that the necessary property of scale-invariance is satisfied (p. 1556, criterion 6, Bayarri et al., 2012). This allows for specifying a prior directly on ρ_{ij} (Section 3.3).

3.2. Exhaustive testing

It is important to extend inference beyond two-sided hypothesis testing (Equation 5). In clinical applications, say, post-traumatic stress disorder (PTSD, McNally et al., 2015; Fried et al., 2018; Armour et al., 2017), a researcher may expect symptoms to share a positive association with one another. This prediction reflects a central tenet of network theory (i.e., symptoms are hypothesized to follow a positive manifold, Borsboom et al., 2011) and

10

it also has a long tradition in psychometrics (see references in Borg, 2018). Furthermore, it may be important to assess the evidence between a positive and negative edge (i.e., "negative manifold," Bork et al., 2019). This can be accomplished by quantifying the relative support for $\rho_{ij} = 0$ vs. $\rho_{ij} < 0$ vs. $\rho_{ij} > 0$. We term this approach *exhaustive testing* and further note that it is a novel contribution to the GGM literature.

Our exhaustive approach simultaneously tests the following hypotheses

$$\mathcal{H}_0 : \rho_{ij} = 0 \tag{7}$$

$$\mathcal{H}_1 : \rho_{ij} > 0$$

$$\mathcal{H}_2 : \rho_{ij} < 0,$$

for each separate partial correlation. The Bayes factor for each hypothesis in (7) is computed from the inverse of the Bayes factor given in (6). This corresponds to the evidence in favor of the unrestricted model (i.e., BF_{1u}). To compute the one-sided Bayes factor for $\mathcal{H}_1$ (BF_{10}), evidence between the unrestricted and order-restricted model is first assessed (BF_{1u}) as

$$\mathrm{BF}_{1u} = \frac{\int_0^\infty p(\rho_{ij}|\mathbf{Y}, \mathcal{H}_u)d\rho_{ij}}{\int_0^\infty p(\rho_{ij}|\mathcal{H}_u)d\rho_{ij}} = \frac{\Pr(\rho_{ij} > 0|\mathbf{Y}, \mathcal{H}_u)}{\Pr(\rho_{ij} > 0|\mathcal{H}_u)}, \tag{8}$$

where the numerator and denominator denote the posterior and prior probability that ρ_{ij} is greater than zero under the unrestricted model $\mathcal{H}_u$ (Equation 5, Klugkist et al., 2005). This can be computed by sampling from the prior and posterior distributions, respectively, or by assuming each follows a normal distribution. The latter leads to an analytic solution,

11

which allows for scaling to large matrices. We present an analytic and sampling-based solution below. Due to the transitive property of the Bayes factor, the relative evidence for $\mathcal{H}_1$ against $\mathcal{H}_0$ is then given by $\mathrm{BF}_{10} = \mathrm{BF}_{u0} \cdot \mathrm{BF}_{1u}$. The Bayes factor for a negative edge is obtained in a similar fashion (resulting in BF_{20}). With the three Bayes factors in hand, each compared to $\mathcal{H}_u$, the posterior hypothesis probabilities can then be computed

$$p(\mathcal{H}_1|\mathbf{Y}) = \frac{p(\mathcal{H}_1)\mathrm{BF}_{1u}}{p(\mathcal{H}_0)\mathrm{BF}_{0u} + p(\mathcal{H}_1)\mathrm{BF}_{1u} + p(\mathcal{H}_2)\mathrm{BF}_{2u}}, \tag{9}$$

where $p(.)$ is the prior probability for each hypothesis. This results in the posterior probability for a positive edge (i.e., $\mathcal{H}_1$). Together, the formulation in (7) provides an *exhaustive* assessment for all individual relations in the network. The ratio of these posterior probabilities can then be used to compare theoretically relevant hypotheses. For example, $\mathrm{BF}_{12} = p(\mathcal{H}_1|\mathbf{Y})/p(\mathcal{H}_2|\mathbf{Y})$ gives the posterior odds in favour of $\rho_{ij} > 0$ compared to $\rho_{ij} < 0$.

3.3. Prior specification

For Bayes factor hypothesis testing, it is important that the encompassing prior for the partial correlations has the following two properties:

I. The prior should be computationally feasible which implies that it should be conjugate or conditionally conjugate.

II. The prior distribution should allow researchers to incorporate the expected edge size (nonzero partial correlations) based on external prior knowledge (e.g., from the psychological network literature).

12

We discuss these properties for two prior distributions: (i) the Wishart prior; and (ii) the matrix$-F$ prior (Mulder and Pericchi, 2018). The latter is novel in the literature on Gaussian graphical modeling.

3.3.1. Wishart prior

The conjugate Wishart prior is the gold standard for modeling precision matrices. Due to its conjugacy, Bayesian computation can be relatively efficient, as will be discussed later, and therefore the prior property I is met. To our knowledge, however, the implied marginal prior for the partial correlations based on a Wishart prior for the precision matrix has not been carefully investigated. This stands in contrast to the inverse Wishart prior for the covariance matrix, as the marginal prior for bivariate correlations is well understood (pp. 1285 - 1287, Barnard et al., 2000).

The Wishart prior under the unconstrained model in (4) is defined by

$$\boldsymbol{\Theta}|\mathcal{H}_u \sim W(\nu_0, \mathbf{B}_0), \tag{10}$$

with $\nu_0 > p - 1$ degree of freedom and a positive definite scale matrix $\mathbf{B}_0$. The covariance matrix $\Sigma \sim IW(\nu_0, \mathbf{B}_0^{-1})$ then follows an inverse-Wishart (IW) distribution. For the following, we set $\nu_0 = p$ and $\mathbf{B}_0 = \mathbf{I}_p$ (i.e., an identity scale matrix), resulting in the standard Wishart distribution (Drton et al., 2008). By setting the degrees of freedom equal to the smallest possible integer, the prior can be said to be minimally informative. Furthermore, the identity scale matrix implies that the off-diagonal elements of the precision matrix are centered around zero. As partial correlations can be seen as normalized versions of the off-diagonal matrix (see Equation 1), this suggests that the implied priors for the partial

correlations are also centered around the null value. Note that centering the prior around the null value already dates back to Jeffreys (1961). When testing partial correlations, this implies that negative and positive partial correlations are equally likely (which seems to be an objective choice), and that small partial correlations are more likely than large partial correlation (Table 2, Wysocki and Rhemtulla, 2019).

Although the marginal prior for ρ_{ij} is not available in closed form, the moments of the Wishart are readily available (e.g. Drton et al., 2008). As a result, we investigated the implied prior by working with known properties of the Wishart distribution and numerically by drawing samples from it. Note that the former leads to an analytic solution for computing the Bayes factor in (6). For brevity, this is provided in the Appendix. For the latter, we sampled from the Wishart distribution with $\nu = p$. Note that this reflects the dimensions, $p \times p$, of a given network. The implied prior was then obtained by standardizing each draw with (A.8). The corresponding densities for $p \in \{3, 20, 50, 100, \text{ and } 200\}$ are visualized in Figure 1. It can be seen that the density is a function of p: the prior is approximately uniform with $p = 3$, whereas, with $p = 200$, the density is narrowly concentrated around zero and there is essentially zero density greater than 0.20 (this often considered a medium effect size, Cohen, 1992). In fact, it can be shown that the marginal prior for a partial correlation is approximately normally distributed with $\mathcal{N}(0, 1/p)$ (see Appendix A.1).

The dependency on p could be reasonable as partial correlations may on average be smaller for larger networks. It is unlikely however that this happens at the exact rate implied by the Wishart prior. Furthermore, the prior would not allow researchers to specify the anticipated magnitude of nonzero partial correlations based on external knowledge
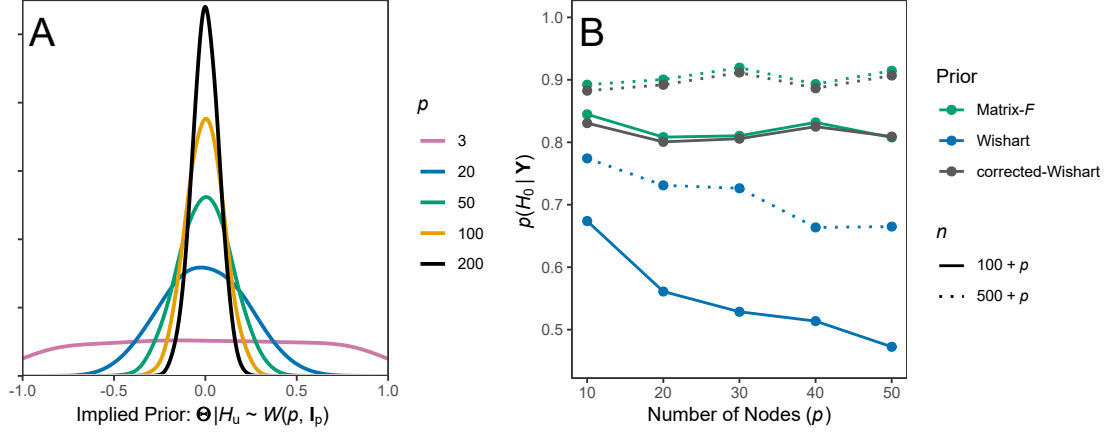
Figure 1: Panel A includes the implied prior for the partial correlations under the Wishart prior. The degrees of freedom, ν (Equation 10), was set to the number of variables, p, in the GGM. This highlights the relation between the prior density and the matrix dimensions. Panel B compares $p(\mathcal{H}_0|\mathbf{Y})$ for the matrix$-F$ (Equation 17), Wishart (Equation 10), and corrected-Wishart (Equation 12) priors. This indicates that p influences the posterior probabilities for the Wishart prior, whereas the corrected-Wishart overcomes this limitation and approximates the matrix$-F$ prior.

(e.g., Table 2 in Wysocki and Rhemtulla, 2019). Therefore, in general the Wishart prior does not satisfy prior property II.

Before introducing an alternative to the Wishart prior which does satisfy both prior properties, it is important to note that the undesirable property of the Wishart prior can be corrected via the prior probabilities of the hypotheses. To see this, note that the posterior odds in (6) are given by the product of the Bayes factor and the prior odds, i.e.,

$$\frac{p(\mathcal{H}_0|\mathbf{Y})}{p(\mathcal{H}_u|\mathbf{Y})} = \frac{p(\rho_{ij} = 0|\mathbf{Y}, \mathcal{H}_u)}{p(\rho_{ij} = 0|\mathcal{H}_u)} \times \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_u)}. \tag{11}$$

The marginal posterior and prior at $\rho_{ij} = 0$ in the Bayes factor can be approximated by simple analytic solutions given their approximate normal distributions (see Appendix A). Now assume that equal prior probabilities are assigned to $\mathcal{H}_0$ and $\mathcal{H}_u$ (i.e., $p(\mathcal{H}_0) =$

$p(\mathcal{H}_u) = 0.5$), and a uniform prior is assumed for the partial correlation under $\mathcal{H}_u$ (so that $p(\rho_{ij} = 0|\mathcal{H}_u) = 0.5$). In that case, the posterior odds in (11) would be computed as

$$\frac{p(\mathcal{H}_0|\mathbf{Y})}{p(\mathcal{H}_u|\mathbf{Y})} = \frac{p(\rho_{ij} = 0|\mathbf{Y}, \mathcal{H}_u)}{0.5} \times \frac{0.5}{0.5}. \tag{12}$$

Now when using the Wishart prior, the approximately normal marginal prior for the partial correlation, $\mathcal{N}(0, 1/p)$, results in a denominator in the first term on the right hand side of (11) of $\sqrt{p/2\pi}$. Thus, if we want to obtain posterior odds that are based on the assumed priors mentioned above while using the Bayes factor based on the Wishart prior in (12), the prior odds should be set to $\frac{p(\mathcal{H}_0)}{p(\mathcal{H}_u)} = \sqrt{2p/\pi}$ in the second term of (11), or in general,

$$\frac{p(\mathcal{H}_0|\mathbf{Y})}{p(\mathcal{H}_u|\mathbf{Y})} = B_{0u} \times \frac{p^W(\rho_{ij} = 0|\mathcal{H}_u)}{0.5} \frac{p(\mathcal{H}_0)}{p(\mathcal{H}_u)}. \tag{13}$$

This method would provide researchers with a fast method to perform exploratory Bayesian testing of partial correlations in GGMs while correcting the Wishart prior via the prior odds for the hypotheses. This is particularly useful because GGMs can include hundreds of nodes (i.e., 1000's of partial correlations). Note also that the denominator in the first term of (12) can accommodate additional priors (i.e., not only a uniform distribution). This would require replacing 0.5 with $p(\rho_{ij} = 0|\mathcal{H}_u)$, given the choice of $\mathcal{H}_u$. The performance of this approximated method will be investigated at the end of this section.

3.3.2. Matrix−F prior

Recently, Mulder and Pericchi (2018) proposed the matrix−F prior distribution as a more flexible alternative to the Wishart and inverse Wishart prior for covariance and precision matrices (for another application of F distributed priors, see, for example, Mulder and Fox (2018)). To our knowledge the matrix−F prior has not yet been considered for modeling the precision matrix of GGMs. We specify an encompassing matrix−F prior distribution for the precision matrix,

$$\mathbf{\Theta} \sim F(\nu, \delta, \mathbf{B}), \tag{14}$$

where $\nu > p - 1$ and $\delta > 0$ are the first and second degrees of freedom, which controls the behavior near the origin and in the tails, respectively, and $\mathbf{B}$ is a positive definite scale matrix. For completeness, the prior density of the matrix−F prior is given in Appendix B.

The matrix−F prior satisfies the prior property I because it can be written as a scale mixture of Wishart distributions with an inverse Wishart mixture distribution, i.e.,

$$\begin{aligned} \mathbf{\Theta}|\mathbf{\Psi} &\sim W(\nu, \mathbf{\Psi}) \\ \mathbf{\Psi} &\sim IW(\delta + p - 1, \mathbf{B}). \end{aligned} \tag{15}$$

Because the Wishart prior is conjugate, the matrix F prior is conditionally conjugate where the conditional posterior of $\mathbf{\Theta}$ given $\mathbf{\Psi}$ has a Wishart distribution, and the conditional posterior of $\mathbf{\Psi}$ given $\mathbf{\Theta}$ has an inverse Wishart distribution (see the collapsed Gibbs sampler in Appendix B where $\boldsymbol{\mu}$ is integrated out). This makes the matrix−F prior computationally feasible for GGMs.

To see whether the matrix−F prior satisfies prior property II, we make use of the fact

17

that the prior can be written as a scale mixture of inverse Wishart distributions with a Wishart mixture distribution, i.e.,

$$\boldsymbol{\Theta}|\boldsymbol{\Phi} \quad \sim \quad IW(\delta + p - 1, \boldsymbol{\Phi}) \tag{16}$$

$$\boldsymbol{\Phi} \quad \sim \quad W(\nu, \mathbf{B}).$$

Furthermore, due to Barnard et al. (2000) it is known that a covariance matrix having an inverse Wishart prior distribution with an identity scale matrix, i.e., $IW(\nu, \mathbf{I}_p)$, results in marginal priors for the bivariate correlations having scaled $beta(\frac{\nu-p+1}{2}, \frac{\nu-p+1}{2})$ distributions on the interval $(-1, 1)$. Consequently, if a precision matrix would have an inverse Wishart prior distribution, i.e., $\boldsymbol{\Theta} \sim IW(\delta + p - 1, \mathbf{I}_p)$, the partial correlations would follow a scaled $beta(\frac{\delta}{2}, \frac{\delta}{2})$ distribution in the interval $(-1, 1)$, which is invariant for the dimension of the network p. We therefore set $\mathbf{B} = \epsilon\mathbf{I}_p$ and $\nu = \epsilon^{-1}$, for a very small value for ϵ, say, 0.00001, so that $\boldsymbol{\Phi} \approx \mathbf{I}_p$ and $\boldsymbol{\Theta}$ is approximately distributed as $IW(\delta + p - 1, \mathbf{I}_p)$.

In sum, the prior for the precision matrix and the implied marginal prior for the partial correlations are specified as

$$\boldsymbol{\Theta} \quad \sim \quad F(\epsilon^{-1}, \delta, \epsilon\mathbf{I}_p) \tag{17}$$

$$\rho_{ij} \quad \sim \quad beta(\tfrac{\delta}{2}, \tfrac{\delta}{2}) \text{ on } (-1, 1),$$

for $i \neq j \in \{1, \ldots, p\}$, respectively. The prior hyperparameter, δ, can be chosen such that the prior standard deviation of ρ_{ij} is approximately equal to the expected deviation of a partial correlation from zero under the unconstrained model. This is given by $\sqrt{\frac{1}{4}(\delta + 1)^{-1}} \cdot 2$, which is the standard deviation of a beta distribution that is multiplied by 2 (due to the

interval length of 2).

Furthermore, Bayesian computation is relatively straightforward as the matrix$-F$ prior is conditionally conjugate for the precision matrix. Thus a posterior sample can straightforwardly be obtained using a Gibbs sampler (see Appendix B). Given the flexibility to tune the prior for partial correlations and its conditional conjugacy, the matrix$-F$ prior both satisfies properties I and II.

As the Bayes factor for a null hypothesis $\mathcal{H}_0 : \rho_{ij} = 0$ versus $\mathcal{H}_u : \rho_{ij} \in (-1, 1)$ is based on the Savage-Dickey density ratio in (6), it is well-known (e.g., Dickey, 1971; Verdinelli and Wasserman, 1995; Heck, 2018) that the prior for the nuisance parameters under $\mathcal{H}_0$, denoted by $\boldsymbol{\phi}_{ij}$ (i.e., all model parameters excluding ρ_{ij}) is implicitly specified as the unconstrained prior of the nuisance parameters under $\mathcal{H}_u$ given $\rho_{ij} = 0$, i.e.,

$$p(\boldsymbol{\phi}_{ij}|\mathcal{H}_0) = p(\boldsymbol{\phi}_{ij}|\mathcal{H}_u, \rho_{ij} = 0) \propto p(\rho_{ij} = 0, \boldsymbol{\phi}_{ij}|\mathcal{H}_u) \tag{18}$$

As the unconstrained prior follows a matrix$-F$ prior, the prior for the nuisance parameters follows a matrix$-F$ distribution under a graph without an edge between nodes i and j. This can be viewed as the matrix$-F$ equivalent of the generalized Wishart distribution for a given graph structure G. The generalized Wishart distribution is proportional to the Wishart distribution when the partial correlations lie in the graph and zero elsewhere (Roverato, 2002; Letac et al., 2007). Given the proven effectiveness of the generalized Wishart for GGMs (e.g., Mohammadi and Wit, 2015), the generalized version of the matrix$-F$ prior seems a reasonable prior for the nuisance parameters under $\mathcal{H}_0$.

19

3.4. Numerical experiments: exploratory

3.4.1. Posterior probabilities for conditional independence

In this experiment, we investigated the performance of the posterior hypothesis probabilities obtained from the Wishart prior, its correction, and the matrix$-F$ prior. To this end, we generated multivariate normal data with an empty network (i.e., no edges), $\mathbf{y} \sim \mathcal{N}_p(\mathbf{0}, \mathbf{I}_p)$, and quantified $p(\mathcal{H}_0|\mathbf{Y})$ for the partial correlation $\rho_{12} = 0$. Note that with more control variables (i.e., a larger network) the posterior standard deviation of ρ_{12} necessarily becomes larger. We accounted for this by setting the sample size to $n + p$, for $n = \{100, 500\}$ and $p = \{10, 20, \ldots, 50\}$. This isolates the effect of the prior density on $p(\mathcal{H}_0|\mathbf{Y})$. The results were averaged across 100 simulation trials.

Figure 1 (panel B) includes these results. The non-invariance of the Wishart prior to p is readily apparent, in that $p(\mathcal{H}_0|\mathbf{Y})$ steadily decreases as the network included more nodes. This is most striking for the smallest sample sizes and $p = 50$, where, on average, $p(\mathcal{H}_0|\mathbf{Y})$ was below 0.50. This indicates that there was more support for $\mathcal{H}_u$, although $\mathcal{H}_0$ was *true*. Our proposed adjustment to the Wishart prior (Equation 12) corrected this behavior and closely mirrored the matrix$-F$ with a uniform prior for ρ_{ij}. In other words, the judicious use of prior hypothesis probabilities can be used to overcome the limitations of using a Wishart prior. Note that for moderate p, say, less than 50, the matrix$-F$ prior only takes a couple seconds to draw 10,000 samples from the posterior. Hence, our analytic solution is ideal for large networks or when an almost instant estimate is desired.

3.4.2. Bayes factor consistency

A desirable feature of Bayes factor is consistency, that is, as n goes to infinity it approaches infinity in favor of the true hypothesis. To investigate Bayes factor consistency,

we quantified $p(\mathcal{H}_0|\mathbf{Y})$ and $p(\mathcal{H}_u|\mathbf{Y})$ in (4). Assuming that, say, $\mathcal{H}_0$ is true, a consistent Bayes factor translates into $p(\mathcal{H}_0|\mathbf{Y})$ approaching one. To this end, we defined the following true precision matrix

$$\mathbf{\Theta} = \begin{bmatrix} 1 & & & & \\ -0.10 & 1 & & & \\ 0 & -0.10 & 1 & & \\ -0.10 & 0 & 0 & 1 & \\ 0 & -0.10 & 0 & -0.10 & 1 \end{bmatrix},$$

which corresponds to partial correlations of 0 and 0.10 (see Equation 1). We then computed the average hypothesis probabilities across 100 trials for n ranging in size from 250 to 100,000. We included both the matrix$-F$ ($\delta = 2$) and the corrected-Wishart priors. Figure E.4 includes these results. It can be seen that the probabilities for each hypothesis approached one, although this happened more quickly for $\mathcal{H}_u$ than $\mathcal{H}_0$. This behavior is reviewed in Johnson and Rossell (2010) and note that it is possible to "gather evidence more quickly" for $\mathcal{H}_0$ by defining non-overlapping hypotheses (Mulder and Fox, 2018; Dablander et al., 2020; Morey and Rouder, 2011). Together, this numerically indicates that both priors provide a consistent Bayes factor for conditional independence and dependence.

3.4.3. Model selection performance

This simulation study addressed three aims: (1) the proposed methods were compared to both an alternative Bayesian method that employs the $G-$Wishart prior in the R package **BDgraph** (inclusion probabilities are used to determine the graph, Mohammadi and Wit, 2015) and significance testing of Fisher-Z partial correlations (NHST, Williams and

Rast, 2019). The latter serves as an important baseline, because there is a defined false positive rate (set to $\alpha = 0.01$); (2) for the proposed methods, performance was evaluated for detecting a *true* null hypothesis of conditional independence; and (3) for matrix$-F$ prior, performance for delta values corresponding to standard deviations of 0.20 ($\delta = 24$) and 0.40 ($\delta \approx 5$) were investigated. We followed the simulation procedure outlined in Epskamp (2016) and Williams et al. (2019). Partial correlations were first estimated from 20 PTSD symptoms (Armour et al., 2017) and then absolute values less than 0.05 were set to zero. This was the *true* network structure for the current simulation. For our methods, we used an arbitrary Bayes factor threshold of 3 (although this is commonly considered moderate or "positive" evidence Kass and Raftery, 1995). These steps were followed for each of 500 simulation trials:

1. Generate multivariate normal data with a given sample size $n \in \{250, 500, 1,000,$ and $10,000\}$.

2. Estimate the conditional dependence $\mathbf{A}^{CD}$ and independence $\mathbf{A}^{CI}$ (only for our methods) structures.

3. Compute two performance measures: sensitivity $\frac{\#TP}{\#TP+\#FN}$ and specificity $\frac{\#TN}{\#TN+\#FP}$. Here TP and FP denote the number of false positives, whereas TN and FN are the number of true and false negatives. For estimating $\mathbf{A}^{CD}$, note that on average specificity is equal to $1 - \alpha$ (Figure 1 in Williams and Rast, 2019) and sensitivity is equal to the average power, $1 - \beta$, across all edges (see Equation 7 in Williams, 2020). $\mathbf{A}^{CI}$ was estimated by reversing the labels (0 and 1) in the *true* network.

Table 1 includes these results. We first discuss the results for estimating $\mathbf{A}^{CD}$. As a point of reference, specificity for NHST was perfectly calibrated to 0.99 (i.e., $1 - \alpha$).

Table 1: Simulation results based on 500 generated datasets (Section 3.4.3).

		Conditional Dependence ($\mathbf{A}^{CD}$)							
		Specificity				Sensitivity			
Method	n	250	500	1,000	10,000	250	500	1,000	10,000
Matrix$-F$ ($\delta \approx 24$)		0.95	0.97	0.98	1	0.41	0.54	0.68	1
Matrix$-F$ ($\delta \approx 5$)		0.98	0.99	0.99	1	0.31	0.45	0.61	1
corrected-Wishart		0.99	1	1	1	0.25	0.39	0.57	1
Wishart		0.97	0.98	0.99	1	0.37	0.51	0.67	1
$G-$Wishart		0.96	0.98	0.99	1	0.39	0.52	0.65	1
NHST		0.99	0.99	0.99	0.99	0.27	0.45	0.64	1
		Conditional Independence ($\mathbf{A}^{CI}$)							
		Specificity				Sensitivity			
Method	n	250	500	1,000	10,000	250	500	1,000	10,000
Matrix$-F$ ($\delta \approx 24$)		0.87	0.85	0.90	1	0.37	0.65	0.79	0.95
Matrix$-F$ ($\delta \approx 5$)		0.62	0.70	0.81	1	0.81	0.88	0.93	0.98
corrected-Wishart		0.50	0.62	0.75	1	0.91	0.94	0.96	0.99
Wishart		0.86	0.83	0.89	1	0.39	0.68	0.81	0.96

This resulted in relatively low sensitivity that was comparable to the corrected-Wishart. Interestingly, the uncorrected Wishart prior was almost identical the $G-$Wishart prior, where the scores were either the same or within 0.02. The advantage of the matrix-F is also apparent. For example, using a large δ value (i.e., a smaller standard deviation) had the highest sensitivity while not sacrificing much in terms of specificity. For the conditional independence structure, $\mathbf{A}^{CI}$, the Wishart and matrix$-F$ prior ($\delta = 24$) had similar performance. This is due to the Wishart prior having a similar prior standard deviation ($1/\sqrt{20} \approx 0.22$ vs. 0.20). Specificity was relatively low for both the corrected-Wishart and the matrix$-F$ prior with $\delta = 5$. This is due to placing too much density on large val-

ues, which, in turn, can erroneously favor the null hypothesis with small sample sizes in particular. Given enough data, however, all prior specifications detected nearly all the null effects with perfect specificity.

There is one important distinction between the Bayesian methods and NHST: the Bayesian methods converged on the *true* model with increasing *n*, that is, they are model selection consistent (Casella et al., 2009). This can be seen by noting the scores for both sensitivity and specificity converged to 1 as the sample size increased, whereas this is not the case for frequentist methods in general (but see Bauer et al., 1988). Importantly, this indicates that our proposed method satisfies the desirable properties of scale-invariance, as well as Bayes factor and model selection consistency (see references in Chib and Kuffner, 2016).

At this point it worth discussing the notion of consistency in reference to our testing strategy, wherein $\frac{1}{2}p(p-1)$ individual hypotheses are tested against the full or unconstrained model. A detailed proof is provided in Leday and Richardson (see Lemma 2, 2018) that demonstrates individual testing is consistent for model selection. For intuition why this works, consider Q partial correlations, that is, (ρ_1, ρ_2, ρ_Q), and we want find those that are zero and nonzero. If we would consider a test of whether one specific ρ, say ρ_q, is equal to zero ($\mathcal{H}_{0q}$) or nonzero ($\mathcal{H}_{1q}$) then, as the sample size grows, the evidence would grow to infinity in favor of $\mathcal{H}_{0q}$ or $\mathcal{H}_{1q}$ if the true values equals 0 or not, respectively, regardless of the true values of the other partial correlations. This holds because the Bayes factor is consistent. Now this will be the case for all ρ_q, for $q = 1, ..., Q$, and therefore our method will be consistent for finding the true zeros and nonzeroes in the vector (ρ_1, ρ_2, ρ_Q) when the sample size grows.

Table 2: Posterior hypothesis probabilities based on 500 generated datasets (Section 3.4.4).

Method	ρ_{ij}	$p(\mathcal{H}_u\|\mathbf{Y})$					
		(0, 0.1]	(0.1, 0.2]	(0.2, 0.3]	(0.3, 0.4]	(0.5, 0.6]	(0.6, 0.7]
matrix$-F$ ($\delta = 24$)		0.58	0.67	0.76	0.89	0.99	1
matrix$-F$ ($\delta \approx 5$)		0.40	0.48	0.61	0.80	0.97	0.99
matrix$-F$ ($\delta = 2$)		0.29	0.37	0.50	0.72	0.94	0.98
corrected-Wishart		0.28	0.34	0.45	0.64	0.90	0.96
$G-$Wishart		0.22	0.33	0.48	0.80	0.98	1

3.4.4. Comparison of posterior probabilities

The matrix$-F$ prior confers advantages for Gaussian graphical modeling. First, placing a prior on the partial correlations (i.e., an effect size) is desirable and not possible with alternative methods. Second, the increased flexibility due to the hyperparameter δ can be harnessed to overcome the informativeness of the Wishart prior (Figure 1, panel A), which, in turn, can provide more support for the true hypothesis. To demonstrate this important aspect of the matrix$-F$ prior, we generated a true network ($p = 20$) where the probability of an edge between two nodes was 10%. The partial correlations were sampled from $uniform(0.05, 0.70)$ and n was set to 50. We then quantified $p(\mathcal{H}_u|\mathbf{Y})$ in (4) for the matrix$-F$ prior, with $\delta = \{24, 5, 2\}$, as well as the corrected-Wishart and $G-$Wishart priors. Recall that he latter provides posterior inclusion probabilities (PIP) computed from a binary indicator. The posterior probabilities were average across 500 simulation trials.

Table 2 includes these results. Although Table 1 indicated that the matrix$-F$ with $\delta = 24$ and the $G-$Wishart provided comparable performance, the posterior probabilities were much larger for the matrix$-F$ prior in Table 2. In fact, across all δ values $p(\mathcal{H}_u|\mathbf{Y})$ was larger for the matrix$-F$ prior. Note that the $G-$Wishart approach employs the median

probability model strategy for selecting the graph (i.e., PIP > 0.50, Barbieri and Berger, 2004) and not the Bayes factor. The matrix$-F$ also had clear advantages compared to the corrected-Wishart prior. For example, while the correction in (13) approximates $\delta = 2$, the posterior probabilities were larger for the matrix$-F$. This can also be seen in Figure 1 (panel A) and Figure E.4. Finally, note that our exhaustive approach would provide the posterior probability for a positive partial correlation in (7). Given the true network included only positive relations (a "positive manifold"), the posterior probabilities would be even larger for this one-sided hypothesis than for two-sided hypotheses considered here.

3.4.5. *Summary*

We emphasize that the similarity between our method and the $G-$Wishart approach for model selection is not too surprising, given these dimensions of data do not pose estimation issues (e.g., compared to $n < p$). However, there are several novel and useful aspects of our proposed methodology:

1. The prior distribution can be specified for the partial correlations (e.g., Equation 17), that is, in reference to a standardized effect size.

2. As a result of point 1, this opened the door for our exhaustive approach in (7). This is especially relevant for psychological networks, because edges are often expected to be positive.

3. The matrix$-F$ prior provides increased flexibility compared to the Wishart prior. This more readily allows for incorporating prior knowledge about the effect size and it can also provide larger posterior hypothesis probabilities (Table 2).

4. Our analytic solution provides a computationally efficient and accurate solution (Figure 1, panel B). This side-steps the cumbersome nature of Bayesian methods

26

for estimating *large* networks (i.e., many nodes) or when an almost instant solution is desired.

3.5. *Empirical application: exploratory*

In all applications, we use PTSD datasets that include 16 symptoms from three proposed clusters (see Table D.4). Although networks would typically include all items, this presents challenges for highlighting important aspects of our methodology and also the example you get unwieldy. Hence, we use only two items from each cluster (see Table D.4), resulting in 15 partial correlations in total ($n = 956$ and $p = 6$). This is "Sample 4" in Fried et al. (see Table 1 therein, 2018). In this example, we adopt the customary perspective of estimating the network and describing the results. In Confirmatory hypothesis testing of order constraints, we demonstrate how exploratory results can then be tested in an independent dataset. This is particularly important because it connects our exploratory and confirmatory approaches. We used the matrix$-F$ prior with $\delta = 24$ (see Table 2 in Wysocki and Rhemtulla, 2019) and drew 50,000 samples from both the posterior and prior distributions.

3.5.1. *Exhaustive testing*

We implemented the exhaustive strategy described in Section 3.2. A posterior hypothesis probability greater than 0.75 was considered evidence for a given hypothesis (i.e., BF > 3 against the complement). Note that a threshold of 3 is arbitrary. In practical applications, this should be decided on with care and a robustness check can be used in the case of prior uncertainty. The conditional dependence structure, $\mathbf{A}^{CD}$ (Figure 2), indicates that all edges for which the Bayes factor indicated a positive partial correlation. Recall that net-
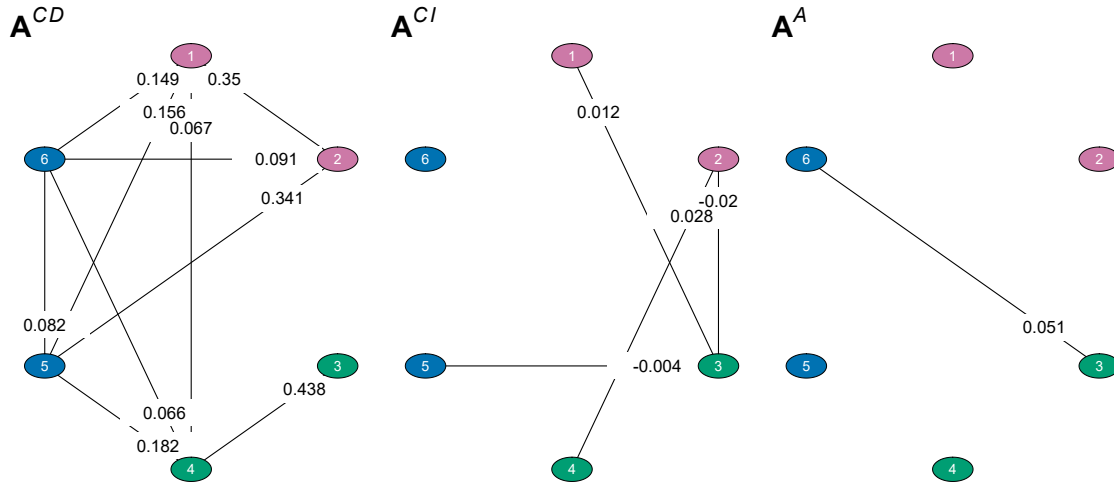
27

Figure 2: Estimated graphical structures. $\mathbf{A}^{CD}$ (Equation 2) is the conditional dependence structure that includes relations supported by the data. This was determined with the exhaustive approach in (7), with $p(\mathcal{H}_1|\mathbf{Y})$ or $p(\mathcal{H}_2|\mathbf{Y})$ greater than 0.75. This corresponds to a Bayes factor greater than three (against the respective complement). $\mathbf{A}^{CI}$ (Equation 3) is the conditional independence structure. This was similarly determined as $p(\mathcal{H}_0|\mathbf{Y}) > 0.75$. $\mathbf{A}^A$ is the ambiguous network that includes those relations for which there was not sufficient evidence to be included in either $\mathbf{A}^{CD}$ or $\mathbf{A}^{CI}$. The node descriptions are provided in Table D.4.

Table 3: Bridge symptoms identified in the exploratory analysis (Figure 2). These results were used to formulate and then test a confirmatory hypothesis (see 22). The node descriptions are provided in Table D.4.

Edge	M	SD	$p(\mathcal{H}_0|\mathbf{Y})$	$p(\mathcal{H}_+|\mathbf{Y})$	$p(\mathcal{H}_-|\mathbf{Y})$
1–4	0.067	0.031	0.186	0.800	0.013
1–5	0.156	0.031	0	0.999	0
2–5	0.341	0.028	0	0.999	0
4–5	0.182	0.31	0	0.999	0
1–6	0.159	0.031	0	0.999	0
2–6	0.091	0.031	0.006	0.992	0.002
4–6	0.066	0.032	0.219	0.766	0.015

work theory expects all positive edges, and, in this case, these symptoms appear to follow a positive manifold. There are also conditionally dependent effects between clusters. These are called "bridge" symptoms and an emerging literature focuses on discovering symptoms that link clusters (see references in Jones et al., 2019a). Note that these inferences are much different than bivariate correlations. In this network, the bridging relation between, say, node one ("intrusive thoughts") and six ("irritability"), indicates an effect that persists after controlling for all other symptoms. The posterior hypothesis probabilities for all bridge symptoms are provided in Table 3.

Evaluating the conditional independence structure, $\mathbf{A}^{CI}$, with Bayesian hypothesis testing is a novel contribution of this work. Hence, Figure 2 includes the first insights into conditionally independent PTSD symptoms. Before describing these results, recall that validated psychometric scales are commonly used in network analysis. By construction, they typically include items with large bivariate correlations. On the other hand, there are several pairs of PTSD symptoms that appear to be conditionally independent. Most

notable is node three ("avoidance thoughts") that can be considered independent of all other symptoms in the network, excluding node four ("avoidance of situation") which belongs to the same cluster. This highlights the powerful inferences made possible with our methodology.

The ambiguous network, $\mathbf{A}^A$, naturally arises from considering both $\mathbf{A}^{CD}$ and $\mathbf{A}^{CI}$. There is only one ambiguous relation between node one ("intrusive thoughts") and four ("avoidance of situations"). We emphasize that this will not always be the case. In these data, the sample size is large and the network is small in size which leads to clearer evidence towards either conditional dependence or independence. When analyzing the full dataset ($p = 16$), for example, $\mathbf{A}^A$ included 22 relations. These insights are not possible with the customary GGM approach that only estimates and visualizes $\mathbf{A}^{CD}$.

4. Confirmatory hypothesis testing of order constraints

Psychological theories can be formulated in the language of mathematics, which, in turn, can be expressed as hypotheses with multiple order constraints on the parameters of interest (Hoijtink, 2011). In a GGM, it may be expected that a set of partial correlations are approximately equal to each other (given a maximal upper bound), larger or smaller than another set of partial correlations, or larger or smaller than a constant (typically zero). These kinds of hypotheses can be derived from a guiding theory or an exploratory analysis. For the latter, note that order-constraints for edges can be read directly from Figure 2 and Table 3. Here the focus is not on determining the graph, but rather Bayesian testing of the interrelations among partial correlations.

In matrix notation, such constrained hypotheses can be formulated as

$$\mathcal{H}_t : \mathbf{R}_t \boldsymbol{\rho} > \mathbf{r}_t, \tag{19}$$

for $t = 1, \ldots, T$ competing hypotheses. In (19), $[\mathbf{R}_t|\mathbf{r}_t]$ is a matrix, $\mathbf{r}$ is a vector, and the inequality specifies the constraints under $\mathcal{H}_t$. The subspace under a constrained hypothesis $\mathcal{H}_t$ that satisfies the constraints on $\boldsymbol{\rho}$ will be denoted by Ω_t. In order to compute Bayes factors between these T hypotheses two important aspects needs to be addressed. First, proper priors need to be specified for the free partial correlations under all hypotheses. Second, once the priors have been formulated, marginal likelihoods under the separate hypotheses need to be computed which can be computationally expensive (Kass and Raftery, 1995).

Both aspects become considerably simplified when adopting an encompassing prior approach (Klugkist et al., 2005). The idea is to specify an unconstrained (or encompassing) prior under an unconstrained model with no constraints on the partial correlations (besides the constraints that ensure positive definiteness). The priors under the constrained hypotheses are then truncations of this encompassing prior truncated in the constrained subspaces. Thus, instead of having to formulate T separate priors, only one unconstrained prior needs to be formulated. Furthermore, due to the encompassing prior approach, marginal likelihood computation can be avoided because the Bayes factor of each constrained hypothesis against the unconstrained hypothesis $\mathcal{H}_u$ can be expressed as the ratio of the posterior and prior probability that the order constraints of $\mathcal{H}_t$ hold under the larger unconstrained model,

$$\mathrm{B}_{tu} = \frac{\mathrm{Pr}(\boldsymbol{\rho} \in \Omega_t | \mathbf{Y}, \mathcal{H}_u)}{\mathrm{Pr}(\boldsymbol{\rho} \in \Omega_t | \mathcal{H}_u)}, \tag{20}$$

31

The general expression first appeared in Klugkist et al. (2005). For order hypotheses on partial correlations, the same steps can be done as for order hypotheses on bivariate correlations (Mulder, 2016). Mulder (2014) referred to the posterior probability in the numerator and the prior probability in the denominator as measures of 'relative fit' and 'relative complexity' of $\mathcal{H}_t$ relative to $\mathcal{H}_u$, respectively. To see this, in the case the order constraints are (not) supported by the data, this results in a large (small) posterior probability, which implies a good (bad) fit of the order constraints. Furthermore, in the case of many (few) constraints, the prior probability that the constraints hold under $\mathcal{H}_u$ will be relatively small (large), which implies a very precise/simple (imprecise/complex), hypothesis. Once (20) has been obtained for all constrained hypotheses of interest, the Bayes factors between competing constrained hypotheses can then be computed using the transitivity property of the Bayes factor, e.g., $\mathrm{BF}_{12} = \frac{\mathrm{BF}_{1u}}{\mathrm{BF}_{2u}}$.

Due to the attractive properties of the matrix$-F$ distribution, the encompassing prior for the confirmatory test will also follow a matrix$-F$ distribution. An efficient method for obtaining the prior and posterior probabilities in (20) is presented in Appendix C. This algorithm for computing Bayes factors for the confirmatory test of order hypotheses in GGMs is implemented in the R package **BGGM**.

4.1. Numerical experiment: confirmatory

This simulation study addressed three aims: (1) we evaluated the sample sizes needed to gain evidence for an order-constrained hypothesis, which also investigates the consistency of the posterior hypothesis probabilities; (2) the number of constraints were varied to examine the posterior probabilities for increasingly complex/simple hypotheses; (3) two values for $\delta = 1, 2$ (Equation 17) were included to evaluate the extent to which the prior

influences the posterior probabilities. For the true GGM, we used the partial correlations estimated from the symptoms included in Figure 2. We then formulated the following *true* hypotheses

$$\mathcal{H}_1^1 : \rho_{34} > \rho_{12} > \rho_{45} \tag{21}$$

$$\mathcal{H}_1^2 : \rho_{34} > \rho_{12} > \rho_{45} > \rho_{56}$$

$$\mathcal{H}_1^3 : \rho_{34} > \rho_{12} > \rho_{45} > \rho_{56} > 0.$$

Note that each hypothesis in (21) was compared to the corresponding complement, $\mathcal{H}_2^i$: "not $\mathcal{H}_1^i$", $i = 1, 2, 3$. To this end, we assumed equal prior probabilities (i.e., $\frac{1}{2}$). Each hypothesis was first compared to the unconstrained model, $\mathcal{H}_u$, that results in BF_{1u}^i and BF_{2u}^i. Note that these Bayes factors are not of substantive interest, but they are needed for then computing BF_{12}^i. We considered sample size sizes ranging from 100 to 1,000 (in increments of 100). The posterior hypothesis probabilities, $p(\mathcal{H}_1^i|\mathbf{Y})$, were averaged across 100 simulation trials.

Figure 3 includes these results. The panels correspond to the order-constraints in (21), indicated by their respective complexity (i.e., the prior probability that the constraints holds under $\mathcal{H}_u$). The posterior hypothesis probabilities approached one as the sample size became larger. Gathering evidence for $\mathcal{H}_1$ was influenced by the number of constraints tested. For example, when testing more specific hypotheses, this was rewarded with larger probabilities in favor of $\mathcal{H}_1$. Note that the hyperparameter, δ, value had a minimal effect on the results. This indicates that the corresponding Bayes factors would be robust to
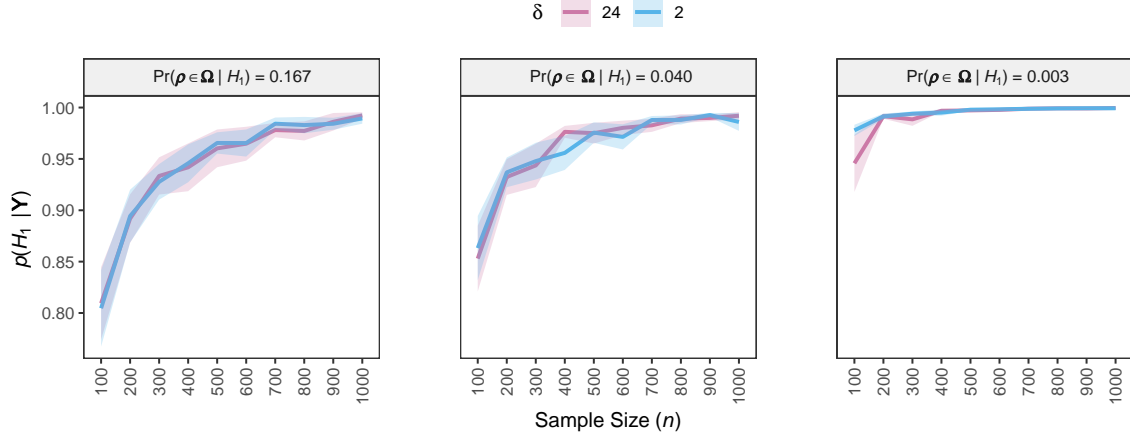
Figure 3: Simulation results (Section 4.1). $p(\mathcal{H}_1|\mathbf{Y})$ is the posterior probability for the hypotheses in (21) compared to the respective complement. $\Pr(\rho \in \Omega|\mathcal{H}_1)$ is the denominator in (20), with smaller values indicating a more specific hypothesis (relative to $\mathcal{H}_u$). The hyperparameter, δ, governs the informativeness of the prior distribution in (17). $\delta = 24$ corresponds to a standard deviation of 0.20 and $\delta = 2$ corresponds to an approximate uniform distribution with a standard deviation of 0.58. The ribbons denote 95 % credible intervals.

different priors (Mulder, 2014).

4.2. Empirical application: confirmatory

In the following, we continue our running example with PTSD symptoms. To this end, we first show how to formulate and test hypotheses from our exploratory method (Figure 2) in an independent dataset ($n = 926$ and $p = 6$). This is "Sample 3" in Fried et al. (2018). The next example considers *a priori*, clinical expectations, regarding clusters of PTSD symptoms. For all models, we employed the matrix$-F$ prior ($\delta = 5$) and drew 50,000 samples from both the posterior and prior distributions. Note that for testing order hypotheses, the Bayes factor is very robust to different prior choices (e.g., the Jeffreys-Lindley paradox is not an issue, Mulder, 2014), and therefore a relatively vague matrix$-F$ prior is used.

34

4.2.1. Bridge symptoms

Identifying symptoms that link clusters in psychopathology networks is an active area of research. Typically, metrics are used to quantify, say, the sum of "bridging" edges for a given node to identify the most central or "important" symptom. In reference to Table 3, this would be accomplished by adding together the partial correlations for node 6 ("irritability"). Conversely, our approach allows for zooming in to examine the interrelations within or between bridging nodes, e.g.,

$$\mathcal{H}_1 : \rho_{25} > \rho_{45} > \rho_{15} > \rho_{35} \tag{22}$$

$$\mathcal{H}_2 : \text{"not } \mathcal{H}_1 \text{."}$$

This hypothesis was formulated from our exploratory testing approach (Table 3) and it captures the order of bridging relations for node 5 ("sleep problems").[2] Further, because nodes 3 ("avoidance of thoughts") and 5 were determined to be conditionally independent (see $\mathbf{A}^{CI}$ in Figure 2), this implies that the relation is smaller than ρ_{15}. In reference to (19), the system of inequalities under $\mathcal{H}_1$ can be formulated as

---

[2]Note that ρ_{35} in Figure 2 is nearly zero which implies that it is less than the other relations. An exact equality constraint, e.g. $\rho_{35} = 0$, is beyond the scope of this work.

$$\mathbf{R}_{\mathcal{H}_1}\boldsymbol{\rho} = \begin{bmatrix} 0 & 1 & 0 & -1 \\ -1 & 0 & 0 & 1 \\ 1 & 0 & -1 & 0 \end{bmatrix} \begin{bmatrix} \rho_{15} \\ \rho_{25} \\ \rho_{35} \\ \rho_{45} \end{bmatrix} > \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}, \tag{23}$$

where $\mathbf{R}_{\mathcal{H}_1}$ denotes a matrix containing the coefficients of the contrasts of interest. This is then tested against it complement, $\mathcal{H}_2$. The Bayes factors were first computed with the unconstrained hypothesis, resulting $\text{BF}_{1u} \approx \frac{0.916}{0.041} = 22.4$ and $\text{BF}_{2u} \approx \frac{0.084}{0.959} = 0.016$[3]. Assuming equal prior model probabilities, $p(\mathcal{H}_1) = p(\mathcal{H}_2) = \frac{1}{2}$, the posterior model probabilities are then $p(\mathcal{H}_1|\mathbf{Y}) > 0.99$ and $p(\mathcal{H}_2|\mathbf{Y}) < 0.01$, which corresponds to overwhelming evidence for $\mathcal{H}_1$ ($\text{BF}_{12} = 256.2$). This indicates that these bridge relations are not only replicable, but there is a specific order among edges bridging the hyper-arousal symptom, "sleep problems," to the other clusters of PTSD symptoms.

4.2.2. Replicating edge order

The topic of replicability has recently captivated the network literature (Forbes et al., 2019; Jones et al., 2019b; Williams, 2020). To assess replicability, it is common to focus on the individual edges with either classical (van Borkulo et al., 2016) or Bayesian testing (Williams et al., 2020). Although the latter has the advantage of directly providing evidence for equality of partial correlations across groups, testing order-constraints allows for asking even more fine grained questions. For example, the exploratory results in Figure 2

---

[3]The denominator also shows the accuracy of our method for computing the prior probabilities ($\frac{1}{4!} = \frac{1}{24} = 0.0417$).

encode ordered hypotheses such as

$$\mathcal{H}_1 : \rho_{34} > (\rho_{12}, \rho_{25}) > (\rho_{15}, \rho_{16}, \rho_{45}) \tag{24}$$

$$\mathcal{H}_2 : \rho_{34} > (\rho_{12}, \rho_{25}) > (\rho_{15}, \rho_{16}, \rho_{45}) > 0$$

$$\mathcal{H}_3 : \text{"not } \mathcal{H}_1 \text{ or } \mathcal{H}_2 \text{."}$$

In (24), ρ_{34} is the relation between nodes 3 ("avoidance of thoughts") and 4 ("avoidance of situations"), which is the largest edge in the exploratory analysis, and the other edges have been grouped according to size. $\mathcal{H}_2$ then adds the further constraint that all edges are positive, that is, they all obey the theoretical expectation of a positive manifold. Note that we have only included partial correlations larger than 0.10, which translates into an ordered hypothesis among the strongest edges in Figure 2. For completeness, we also included $\mathcal{H}_3$ that can be understood as neither $\mathcal{H}_1$ or $\mathcal{H}_2$. The Bayes factors were first computed with the unconstrained hypothesis, resulting in $\text{BF}_{1u} \approx \frac{0.77}{0.0154} = 49.4$, $\text{BF}_{2u} \approx \frac{0.63}{0.0001} = 6300$, and $\text{BF}_{3u} \approx \frac{0.0139}{0.9807} = 0.014$. Assuming equal prior model probabilities, $p(\mathcal{H}_1) = p(\mathcal{H}_2) = p(\mathcal{H}_3) = \frac{1}{3}$, the posterior model probabilities are then $p(\mathcal{H}_1|\mathbf{Y}) \approx 0.01$, $p(\mathcal{H}_2|\mathbf{Y}) \approx 0.99$ and $p(\mathcal{H}_3|\mathbf{Y}) < 0.001$. Whereas there was evidence for $\mathcal{H}_1$ ($\text{BF}_{13} > 3000$) and $\mathcal{H}_2$ ($\text{BF}_{23} > 300000$) compared to $\mathcal{H}_3$, the additional constraint to positive values indicated overwhelming evidence for $\mathcal{H}_1$ ($\text{BF}_{12} = 93.7$). In other words, an order-constrained hypothesis derived from an initial, exploratory analysis, was confirmed in an independent dataset. This implies that there is a replicable order, constrained to positive values, for the strongest edges in these PTSD symptoms.

4.2.3. Clinical expectations

The previous examples demonstrated how our exploratory and confirmatory approaches can work in tandem. In this example, we formulate hypotheses directly from the Diagnostic and Statistical Manual of Mental Disorders edition four (DSM-IV). In the DSM-IV, PTSD are thought to have three clusters (experiencing symptoms, avoidance, and hyper-arousal). This clinical expectation implies that partial correlations within a cluster are larger than those between clusters. In other words, there are implicit order-constrained hypotheses, say,

$$\mathcal{H}_1 : \rho_{56} > (\rho_{15}, \rho_{25}, \rho_{35}, \rho_{45}, \rho_{16}, \rho_{26}, \rho_{36}, \rho_{46}) \qquad (25)$$

$$\mathcal{H}_2 : \text{"not } \mathcal{H}_1\text{,"}$$

which focuses on the hyper-arousal cluster. $\mathcal{H}_1$ expresses that the relation within the cluster, ρ_{56}, is larger than all the relations with the other two clusters. This is again tested against the complement, $\mathcal{H}_2$. The Bayes factors were first computed with the unconstrained hypothesis, resulting in $\text{BF}_{1u} \approx \frac{0.003}{0.111} = 0.026$ and $\text{BF}_{2u} \approx \frac{0.997}{0.889} = 1.122$. Assuming equal prior model probabilities, $p(\mathcal{H}_1) = p(\mathcal{H}_2) = \frac{1}{2}$, the posterior model probabilities are then $p(\mathcal{H}_1|\mathbf{Y}) = 0.023$ and $p(\mathcal{H}_2|\mathbf{Y}) = 0.977$. This results in a Bayes factor providing evidence *against* $\mathcal{H}_1$ ($\text{BF}_{12} = 0.024$ or $\text{BF}_{21} = 42.5$), that is, these data do not support the notion that conditional dependencies are larger within the hyper-arousal cluster.

5. Conclusion

We introduced novel methodology for Bayesian hypothesis testing in Gaussian graphical models. The tests were specifically developed for both exploratory and confirmatory hypothesis testing. The former is the most common approach and our exploratory test is the first to provide a Bayes factor for one-sided hypotheses in GGMs. This is especially relevant for psychological applications because partial correlations are often expected to be positive. We extended the testing possibilities to confirmatory settings with order-constraints on multiple partial correlation coefficients. The application demonstrated that the exploratory and confirmatory approaches can work in tandem: hypotheses can be formulated from an initial analysis and then tested in an independent dataset. Furthermore, the confirmatory approach allows for testing theoretical and clinical expectations in GGMs.

Additionally, we introduced the matrix$-F$ prior distribution for Bayesian testing in GGMs. As we demonstrated, this prior offers advantages compared to the customary Wishart and G-Wishart priors. For example, the matrix$-F$ prior is invariant to the dimensions of the precision matrix and it can be specified directly on the partial correlations (i.e., in reference to an effect size). This increased flexibility translated into providing more support for the true hypothesis (Table 2). Moreover, we also provided a computationally efficient, analytic solution, based upon the Wishart prior. This is ideal for large p settings or when an almost instant estimate is needed. Both the matrix$-F$ and Wishart prior were shown to provide either superior or competitive performance compared to the G-Wishart prior. The methodology is implemented in the R package **BGGM**.

Acknowledgements

References

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, 49–59. doi:10.1016/j.janxdis.2016.11.008.

Baba, K., Shibata, R., Sibuya, M., 2004. Partial Correlation and Conditional Correlation as Measures of Conditional Independence. Australian & New Zealand Journal of Statistics 46, 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, 1–19.

Barbieri, M.M., Berger, J.O., 2004. Optimal predictive model selection. Annals of Statistics 32, 870–897. doi:10.1214/009053604000000238, arXiv:0406464v1.

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

Bauer, P., PÖTSCHER, B.M., Hackl, P., 1988. Model Selection by Multiple Test Procedures. Statistics 19, 39–44. doi:`10.1080/02331888808802068`.

Bayarri, M.J., Berger, J.O., Forte, A., Garc\'\ia-Donato, G., others, 2012. Criteria for Bayesian model choice with application to variable selection. The Annals of statistics 40, 1550–1577.

Bhadra, A., Mallick, B.K., 2013. Joint high-dimensional Bayesian variable and covariance selection with an application to eQTL analysis. Biometrics 69, 447–457. doi:`10.1111/biom.12021`.

Böing-Messing, F., van Assen, M.A.L.M., Hofman, A.D., Hoijtink, H., Mulder, J., 2017. Bayesian evaluation of constrained hypotheses on variances of multiple independent groups. Psychological Methods 22, 262.

Borg, I., 2018. A note on the positive manifold hypothesis. Personality and Individual Differences 134, 13–15.

Bork, R.v., Marsman, M., Rhemtulla, M., Epskamp, S., Kruis, J., Borsboom, D., 2019. Common effect models: Positive or negative manifold? PsyArXiv doi:`10.31234/OSF.IO/XS4CY`.

van Borkulo, C.D., Boschloo, L., Kossakowski, J.J., Tio, P., Schoevers, R.A., Borsboom, D., Waldorp, L.J., 2016. Comparing network structures on three aspects: A permutation test. Manuscript submitted , 34doi:`10.13140/RG.2.2.29455.38569`.

Borsboom, D., Cramer, A.O., 2013. Network Analysis: An Integrative Approach to

the Structure of Psychopathology. Annual Review of Clinical Psychology 9, 91–121. doi:10.1146/annurev-clinpsy-050212-185608.

Borsboom, D., Cramer, A.O.J., Schmittmann, V.D., Epskamp, S., Waldorp, L.J., 2011. The Small World of Psychopathology. PLoS ONE 6, e27407. doi:10.1371/journal.pone.0027407.

Braeken, J., Mulder, J., Wood, S., 2015. Relative Effects at Work: Bayes Factors for Order Hypotheses. Journal of Management 41, 544–573. doi:10.1177/0149206314525206.

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

Chib, S., Kuffner, T.A., 2016. Bayes factor consistency. arXiv .

Cohen, J., 1992. A power primer. Psychological Bulletin 112, 155–159. doi:10.1037/0033-2909.112.1.155.

Costantini, G., Epskamp, S., Borsboom, D., Perugini, M., Mõttus, R., Waldorp, L.J., Cramer, A.O., 2015. State of the aRt personality research: A tutorial on network analysis of personality data in R. Journal of Research in Personality 54, 13–29. doi:10.1016/j.jrp.2014.07.003.

Cramer, A.O.J., Waldorp, L.J., van der Maas, H.L.J., Borsboom, D., 2010. Comorbidity: A network perspective. Behavioral and Brain Sciences 33, 137–150. doi:10.1017/S0140525X09991567.

Dablander, F., Bergh, D.v.d., Ly, A., Wagenmakers, E.J., 2020. Default Bayes Factors for Testing the (In)equality of Several Population Variances. arXiv .

Dablander, F., Hinne, M., 2019. Node centrality measures are a poor substitute for causal inference. Scientific reports 9, 1–13. doi:10.1038/s41598-019-43033-9.

D'Angelo, G.M., Luo, J., Xiong, C., 2012. Missing Data Methods for Partial Correlations. Journal of biometrics & biostatistics 3. doi:10.4172/2155-6180.1000155.

Dempster, A., 1972. Covariance Selection. Biometrics 28, 157–175. doi:10.2307/2528966.

Di Pierro, R., Costantini, G., Benzi, I.M.A., Madeddu, F., Preti, E., 2019. Grandiose and entitled, but still fragile: A network analysis of pathological narcissistic traits. Personality and Individual Differences 140, 15–20. doi:10.1016/j.paid.2018.04.003.

Dickey, J.M., 1971. The Weighted Likelihood Ratio, Linear Hypotheses on Normal Location Parameters. The Annals of Mathematical Statistics 42, 204–223. doi:10.1214/aoms/1177693507.

Dittrich, D., Leenders, R., Mulder, J., 2020. Bayesian analysis of higher-order network autocorrelation models. Sociological Methodology .

Dittrich, D., Leenders, R.T.A.J., Mulder, J., 2019. Network Autocorrelation Modeling: A Bayes Factor Approach for Testing (Multiple) Precise and Interval Hypotheses. Sociological Methods & Research 48, 642–676. doi:10.1177/0049124117729712.

Drton, M., Massam, H., Olkin, I., others, 2008. Moments of minors of Wishart matrices. The Annals of Statistics 36, 2261–2283.

Drton, M., Perlman, M.D., 2004. Model selection for Gaussian concentration graphs. Biometrika 91, 591–602. doi:10.1093/biomet/91.3.591.

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

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

Forbes, M.K., Wright, A.G., Markon, K.E., Krueger, R.F., 2019. Quantifying the Reliability and Replicability of Psychopathology Network Characteristics. Multivariate Behavioral Research doi:10.1080/00273171.2019.1616526.

Fried, E.I., Eidhof, M.B., Palic, S., Costantini, G., Dijk, H.M.H.v., Bockting, C.L.H., Engelhard, I., Armour, C., Nielsen, A.B.S., 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, 335–351. doi:10.1177/2167702617745092.

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

Giudici, P., 1995. Bayes factors for zero partial covariances. Journal of Statistical Planning and Inference 46, 161–174. doi:10.1016/0378-3758(94)00101-Z.

Gutiérrez-Peña, E., Smith, A.F.M., Bernardo, J.M., Consonni, G., Veronese, P., George, E.I., Girón, F.J., Martínez, M.L., Letac, G., Morris, C.N., 1997. Exponential and

bayesian conjugate families: Review and extensions. Test 6, 1–90. doi:`10.1007/BF02564426`.

Haslbeck, J.M., Ryan, O., Robinaugh, D., Waldorp, L., Borsboom, D., 2019. Modeling Psychopathology: From Data Models to Formal Theories. PsyArXiv doi:`10.31234/OSF.IO/JGM7F`.

Heck, D.W., 2018. A caveat on the Savage-Dickey density ratio: The case of computing Bayes factors for regression parameters. British Journal of Mathematical and Statistical Psychology doi:`10.1111/bmsp.12150`.

Heck, D.W., Davis-Stober, C.P., 2019. Multinomial models with linear inequality constraints: Overview and improvements of computational methods for Bayesian inference. Journal of mathematical psychology 91, 70–87.

Hoijtink, H., 2011. Informative hypotheses: Theory and practice for behavioral and social scientists. Chapman and Hall/CRC.

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

Horn, J.L., Cattell, R.B., 1966. Refinement and test of the theory of fluid and crystallized general intelligences. Journal of educational psychology 57, 253.

Hsu, C.W., Sinay, M.S., Hsu, J.S., 2012. Bayesian estimation of a covariance matrix with flexible prior specification. Annals of the Institute of Statistical Mathematics 64, 319–342. doi:`10.1007/s10463-010-0314-5`.

Jeffreys, H., 1961. The theory of probability. Oxford University Press, Oxford.

Johnson, V.E., Rossell, D., 2010. On the use of non-local prior densities in Bayesian hypothesis tests. Journal of the Royal Statistical Society: Series B (Statistical Methodology) 72, 143–170. doi:`10.1111/j.1467-9868.2009.00730.x`.

Jones, P.J., Ma, R., McNally, R.J., 2019a. Bridge Centrality: A Network Approach to Understanding Comorbidity. Multivariate Behavioral Research , 1–15doi:`10.1080/00273171.2019.1614898`.

Jones, P.J., Mair, P., McNally, R.J., 2018. Visualizing psychological networks: A tutorial in R. Frontiers in Psychology 9, 1742.

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

Kass, R.E., Raftery, A.E., 1995. Bayes Factors. Journal of the American Statistical Association 90, 773–795.

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

Klugkist, I., Kato, B., Hoijtink, H., 2005. Bayesian model selection using encompassing priors. Statistica Neerlandica 59, 57–69. doi:`10.1111/j.1467-9574.2005.00279.x`.

Krishnamoorthy, K., Xia, Y., 2007. Inferences on correlation coefficients: One sample, independent, and correlated cases. Journal of Statistical Planning and Inference 137, 2362–2379.

Kshirsagar, A.M., 1959. Bartlett Decomposition and Wishart Distribution. The Annals of Mathematical Statistics 30, 239–241. doi:10.1214/AOMS/1177706379.

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, 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, e0148171. doi:10.1371/journal.pone.0148171.

Lauritzen, S.L., 1996. Graphical models. volume 17. Clarendon Press.

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

Leonard, T., Hsu, J.S.J., 1992. Bayesian Inference for a Covariance Matrix Tom Leonard; John S. J. Hsu. Annals of Statistics 20, 1669–1696.

Letac, G., Massam, H., others, 2007. Wishart distributions for decomposable graphs. The Annals of Statistics 35, 1278–1323.

Li, Y., Craig, B.A., Bhadra, A., 2017. The Graphical Horseshoe Estimator for Inverse Covariance Matrices .

Marsman, M., Wagenmakers, E.J., 2017. Bayesian benefits with JASP. European Journal of Developmental Psychology 14, 545–555. doi:10.1080/17405629.2016.1259614.

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

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

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

Morey, R.D., Rouder, J.N., 2011. Bayes Factor Approaches for Testing Interval Null Hypotheses. Psychological Methods 16, 406–419. doi:10.1037/a0024377.

Mulder, J., 2014. Bayes factors for testing inequality constrained hypotheses: Issues with Prior Specification. British Journal of Statistical and Mathematical Psychology 67, 153–171.

Mulder, J., 2016. Bayes factors for testing order-constrained hypotheses on correlations. Journal of Mathematical Psychology 72, 104–115. doi:10.1016/j.jmp.2014.09.004.

Mulder, J., Fox, J.P., 2013. Bayesian tests on components of the compound symmetry covariance matrix. Statistics and Computing 23, 109–122.

Mulder, J., Fox, J.P., 2018. Bayes factor testing of multiple intraclass correlations. Bayesian Analysis .

Mulder, J., Hoijtink, H., de Leeuw, C., 2012. BIEMS: A Fortran 90 program for calculating {B}ayes factors for inequality and equality constrained model. Journal of Statistical Software 46.

Mulder, J., Pericchi, L., 2018. The Matrix-F Prior for Estimating and Testing Covariance Matrices. Bayesian Analysis , 1–22doi:10.1214/17-BA1092.

Mulder, J., Raftery, A.E., 2019. BIC Extensions for Order-constrained Model Selection. Sociological Methods & Research , 004912411988245doi:10.1177/0049124119882459.

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

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

Peng, J., Wang, P., Zhou, N., Zhu, J., 2009. Partial Correlation Estimation by Joint Sparse Regression Models. Journal of the American Statistical Association I 104, 735–746. doi:10.1198/jasa.2009.0126.

Preacher, K.J., 2006. Testing complex correlational hypotheses with structural equation models. Structural Equation Modeling 13, 520–543.

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.

Ryan, O., Bringmann, L., Schuurman, N.K., 2019. The Challenge of Generating Causal Hypotheses Using Network Models doi:10.31234/OSF.IO/RYG69.

Steiger, J.H., 1980. Tests for comparing elements of a correlation matrix. Psychological bulletin 87, 245.

Tsukuma, H., 2014. Bayesian estimation of a bounded precision matrix. Journal of Multivariate Analysis 127, 160–172. doi:10.1016/j.jmva.2014.02.016.

Van Kempen, G.M.P., Van Vliet, L.J., 2000. Mean and variance of ratio estimators used in fluorescence ratio imaging. Cytometry 39, 300–305. doi:10.1002/(SICI)1097-0320(20000401)39:4<300::AID-CYTO8>3.0.CO;2-O.

Verdinelli, I., Wasserman, L., 1995. Computing Bayes Factors Using a Generalization of the Savage-Dickey Density Ratio. Journal of the American Statistical Association 90, 614–618. doi:10.1080/01621459.1995.10476554.

Wagenmakers, E.J., Lodewyckx, T., Kuriyal, H., Grasman, R., 2010. Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. Cognitive Psychology 60, 158–189. doi:10.1016/j.cogpsych.2009.12.001.

Wagenmakers, E.J., Verhagen, J., Ly, A., 2016. How to quantify the evidence for the absence of a correlation. Behavior Research Methods 48, 413–426.

Werner, M., Štulhofer, A., Waldorp, L., Jurin, T., 2018. A Network Approach to Hypersexuality: Insights and Clinical Implications. Journal of Sexual Medicine 15, 410–415. doi:10.1016/j.jsxm.2018.01.009.

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

Williams, D.R., 2018. Bayesian Estimation for Gaussian Graphical Models: Structure Learning, Predictability, and Network Comparisons. PsyArXiv doi:10.31234/osf.io/x8dpr.

Williams, D.R., 2020. Learning to Live with Sampling Variability: Expected Replicability in Partial Correlation Networks. PsyArXiv doi:10.31234/OSF.IO/FB4SA.

Williams, D.R., Piironen, J., Vehtari, A., Rast, P., 2018. Bayesian Estimation of Gaussian Graphical Models with Projection Predictive Selection .

Williams, D.R., Rast, P., 2019. Back to the basics: Rethinking partial correlation network methodology. British Journal of Mathematical and Statistical Psychology doi:10.1111/bmsp.12173.

Williams, D.R., Rast, P., Pericchi, L.R., Mulder, J., 2020. Comparing Gaussian graphical models with the posterior predictive distribution and Bayesian model selection. Psychological Methods .

Williams, D.R., Rhemtulla, M., Wysocki, A.C., Rast, P., 2019. On Nonregularized

Estimation of Psychological Networks. Multivariate Behavioral Research 54, 1–23. doi:10.1080/00273171.2019.1575716.

Wysocki, A.C., Rhemtulla, M., 2019. On Penalty Parameter Selection for Estimating Network Models. Multivariate Behavioral Research , 1–15doi:10.1080/00273171.2019.1672516.

Appendix A. Derivations

Appendix A.1. Implied 1st and 2nd Moments

Our computationally efficient method requires deriving the implied prior variance for the partial correlations. This can be accomplished by working with the known moments of the Wishart distribution. To see this, first consider that we need to solve for the following ratio

$$\text{Var}\left[\frac{w_{ij}}{\sqrt{w_{ii}w_{jj}}}\right] \approx \frac{\text{Var}[w_{ij}]}{\text{E}[w_{ii}] \cdot \text{E}[w_{jj}]}, \tag{A.1}$$

where $\mathbf{W} \sim W_p(p, \mathbf{I}_p)$ has a standard Wishart distribution. We assume that $v = p$ and an identity matrix, $\mathbf{I}_p$, for the scale matrix. Importantly, $\mathbf{W}$ can be decomposed as $\mathbf{W} = \mathbf{L}\mathbf{L}'$, where $\mathbf{L}$ is the lower-triangular matrix and $\mathbf{L}'$ is the transpose. This is the customary Choleski-decomposition, and, in the case of the Wishart, it is known as the Bartlett decomposition that can be used to generate random matrices. The diagonal elements of $\mathbf{L}$, l_{ii}, $i = 1, \ldots, p$, are sampled from a chi-squared distribution $l_{ii}^2 \sim \chi_{v-i+1}^2$. In turn, the off-diagonal elements, l_{ij}, $1 \le j < i \le p$, are drawn from a standard normal distribution $l_{ij} \sim \mathcal{N}(0, 1)$. Note that all of the diagonal and off-diagonal elements are "independently

distributed." We refer interested readers to Kshirsagar (1959, p. 240). This implies that prior distribution for the nuisance parameters in (18) is the same for both the unconstrained and equality constrained models in (5).

To derive (A.1), we work with the known properties of the Wishart distribution. The expectation is given as

$$E[\mathbf{W}] = \nu \mathbf{I}_p, \tag{A.2}$$

which results in the numerator of (A.1) reducing to p^2. This is due to the off-diagonal elements being independent, such that $E[w_{ii} \cdot w_{jj}] = E[w_{ii}] \cdot E[w_{jj}]$. Furthermore, the variance follows

$$\text{Var}[\mathbf{W}_{ij}] = \nu(\mathbf{I}_{ij}^2 + \mathbf{I}_{ii} \cdot \mathbf{I}_{jj}), \tag{A.3}$$

which similarly reduces to p. Together, this makes clear that the implied prior distribution is a function of p, that is,

$$\text{Var}(\rho_{ij}) \approx f(p) \sim \frac{1}{p}, \tag{A.4}$$

with the standard deviation as $\text{SD}(\rho_{ij}) = 1/\sqrt{p}$. When assuming normality, the implied prior distribution for the partial correlations is

$$\rho_{ij} \sim \mathcal{N}(0, \frac{1}{\sqrt{p}}). \tag{A.5}$$

This can also be seen from considering the Taylor series expansion (Equation 6 in Van Kempen and Van Vliet, 2000). First denote the numerator and denominator in (A.1) as a and b, respectively. The variance, assuming that b has support $[0, \infty]$, is then given as

$$\text{Var}\left[\frac{a}{b}\right] \approx \frac{\text{Var}[a]}{\text{E}[b]^2} - \frac{2\text{E}[a]}{\text{E}[b]^3}\text{Cov}[a, b] + \frac{\text{E}[a]}{\text{E}[b]^4}\text{Var}[b]. \tag{A.6}$$

In (A.6), all terms with $\text{E}[a]$ in the numerator cancel out. This is due to having an expectation of zero. This leaves the first term. With, say, $p = 10$, this results in $p/\sqrt{p^4} = p/p^2 = 1/p$. The accuracy of this derivation was verified by sampling from a Wishart distribution and then standardizing the draws with (A.8) for $v = \in \{5, 10, 20\}$. Rounded to three decimal places, the solution based on $1/\sqrt{p}$ and sampling were numerically equivalent to three decimal places: 0.447, 0.316, and 0.223. Furthermore, as seen in Figure 1 (panel A), approximate normality is achieved with $v > 10$.

Appendix A.2. Approximated posterior under the Wishart prior

Given a Wishart prior, the posterior mean has an analytic expression, that is

$$\text{E}[\boldsymbol{\Theta}|\mathbf{Y}] = (n + p)(\mathbf{S} + \mathbf{I}_p)^{-1}, \tag{A.7}$$

where **S** is the scatter matrix, $\mathbf{Y}'\mathbf{Y}$, and the scale matrix has been set to an identity matrix $\mathbf{I}_p$. The $p \times p$ partial correlation matrix, **P**, can be computed as

$$E[\mathbf{P}|\mathbf{Y}] = -\Big([\text{diag}(\boldsymbol{\theta})]^{-1}\boldsymbol{\Theta}[\text{diag}(\boldsymbol{\theta})]^{-1} - \mathbf{I}_p\Big), \tag{A.8}$$

where $\text{diag}(\boldsymbol{\theta})$ is a $p \times p$ diagonal matrix with $\sqrt{\theta_{ii}}$, $i, \ldots, p$, as the diagonal elements. Note that the unique partial correlations $\frac{1}{2}p(p-1)$ are in the lower and upper-triangular of **P**, with the latter denoted as ρ_{ij} for $1 \leq i < j \leq p$. There is a closed form solution for $\text{Var}[\theta_{ij}]$ (Equation A.3), but, in this case, the variances for the partial correlations are required. Because the scale matrix in (10) can be understood as a prior guess for the covariance matrix, $\mathbf{I}_p$ can be considered non-informative. Accordingly, we assume the posterior variance is equivalent to the standard error in likelihood only estimation (Gelman et al., 2014, p. 88). Following D'Angelo et al. (Equation 6 therein, 2012), the approximate variance is

$$\text{Var}[\rho_{ij}|\mathbf{Y}] \approx \frac{1 - r_{ij}^2}{n - c - 2} \tag{A.9}$$

where r_{ij} is the sample based partial correlation and c is the number of variables controlled for (i.e., $p - 2$). Although the Pearson correlation coefficient follows an approximate t−distribution (p. 2 in D'Angelo et al., 2012), we assume that posterior distributions are normally distributed. This simplification is due to GGMs most commonly being fitted with hundreds of observations in the behavioral sciences. The results in Figures 1 (panel

B) and E.4, as well as the results in Tables 1 and 2, demonstrate that this solution provides an accurate approximation to the matrix$-F$ prior.

Appendix B. Gibbs sampler

For a precision matrix $\mathbf{\Theta}$ having a matrix-$F(\nu, \delta, \mathbf{B})$ prior distribution, the probability density function is given by

$$p^F(\mathbf{\Theta}) = \frac{\Gamma_k(\frac{\nu+\delta+p-1}{2})}{\Gamma_k(\frac{\nu}{2})\Gamma_k(\frac{\delta+p-1}{2})|\mathbf{B}|^{\frac{\nu}{2}}}|\mathbf{\Theta}|^{\frac{\nu-p-1}{2}}|\mathbf{I}_k + \mathbf{\Theta}\mathbf{B}^{-1}|^{-\frac{\nu+\delta+p-1}{2}}.$$

Following Mulder and Pericchi (2018), this implies that the covariance matrix $\mathbf{\Sigma}$ follows a matrix-$F(\delta + p - 1, \nu - p + 1, \mathbf{B}^{-1})$, which can be written as a scale mixture of inverse Wishart distributions, i.e.,

$$\mathbf{\Sigma}|\mathbf{\Psi} \sim IW(\nu, \mathbf{\Psi})$$
$$\mathbf{\Psi} \sim W(\delta + p - 1, \mathbf{B}^{-1}).$$

Furthermore, the likelihood for n independent observations from the multivariate normal model is given by

$$p(\mathbf{Y}|\boldsymbol{\mu}, \mathbf{\Sigma}) \propto |\mathbf{\Sigma}|^{-n/2} \exp\left\{-\tfrac{1}{2}\mathrm{tr}\mathbf{\Sigma}^{-1}\mathbf{S}\right\} \exp\left\{-\tfrac{n}{2}(\boldsymbol{\mu} - \bar{\mathbf{y}})'\mathbf{\Sigma}^{-1}(\boldsymbol{\mu} - \bar{\mathbf{y}})\right\},$$

where $\mathbf{S}$ is the sums of squares matrix and $\bar{\mathbf{y}}$ is the vector of sample means. Therefore when using a flat prior for the nuisance parameter $\boldsymbol{\mu}$, the posterior (including the parameter

expansion with $\mathbf{\Psi}$) can be written as

$$
\begin{aligned}
p(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \mathbf{\Psi}|\mathbf{Y}) \quad &\propto \quad p^{IW}(\boldsymbol{\Sigma}|\mathbf{\Psi})p^{W}(\mathbf{\Psi})p(\mathbf{Y}|\boldsymbol{\mu}, \boldsymbol{\Sigma}) \\
&\propto \quad |\mathbf{\Psi}|^{(v+\delta-2)/2}|\boldsymbol{\Sigma}|^{-(n+v+p+1)/2}\exp\left\{-\tfrac{1}{2}\mathrm{tr}\mathbf{\Psi}\boldsymbol{\Sigma}^{-1}\right\} \\
&\qquad \exp\left\{-\tfrac{1}{2}\mathrm{tr}\mathbf{\Psi}\mathbf{B}\right\}\exp\left\{-\tfrac{1}{2}\mathrm{tr}\boldsymbol{\Sigma}^{-1}\mathbf{S}\right\} \\
&\qquad \exp\left\{-\tfrac{n}{2}(\boldsymbol{\mu}-\bar{\mathbf{y}})'\boldsymbol{\Sigma}^{-1}(\boldsymbol{\mu}-\bar{\mathbf{y}})\right\} \\
\Rightarrow p(\boldsymbol{\Sigma}, \mathbf{\Psi}|\mathbf{Y}) \quad &\propto \quad |\mathbf{\Psi}|^{(v+\delta-2)/2}|\boldsymbol{\Sigma}|^{-(n+v+p)/2}\exp\left\{-\tfrac{1}{2}\mathrm{tr}\mathbf{\Psi}\boldsymbol{\Sigma}^{-1}\right\} \\
&\qquad \exp\left\{-\tfrac{1}{2}\mathrm{tr}\mathbf{\Psi}\mathbf{B}\right\}\exp\left\{-\tfrac{1}{2}\mathrm{tr}\boldsymbol{\Sigma}^{-1}\mathbf{S}\right\}.
\end{aligned}
$$

Hence a Gibbs sampler can then be formed using the following conditional posteriors

$$
\begin{aligned}
\boldsymbol{\Sigma}|\mathbf{\Psi}, \mathbf{Y} \quad &\sim \quad IW(n+v-1, \mathbf{S}+\mathbf{\Psi}) \\
\mathbf{\Psi}|\boldsymbol{\Sigma}, \mathbf{Y} \quad &\sim \quad W(v+\delta+p-1, (\boldsymbol{\Sigma}^{-1}+\mathbf{B})^{-1}).
\end{aligned}
$$

A posterior sample for the covariance matrix $\boldsymbol{\Sigma}$, and thus of the precision matrix $\boldsymbol{\Theta}$ and the partial correlations, can be obtained by iteratively sampling $\boldsymbol{\Sigma}$ and $\mathbf{\Psi}$ from their respective conditional posterior distribution.

Appendix C. Computation of prior and posterior probabilities

In order to obtain the Bayes factors in (20), the posterior and prior probability need to be computed. A naive way to compute these probabilities would be to get posterior and prior samples under the unconstrained model, and compute the proportions of draws satisfying the constraints. This method however will be computationally expensive in the case of many constraints which results in very small posterior and prior probabilities. For

example, when the prior assumes that all possible orderings would be (approximately) equally likely, the prior probability of a specific ordering of 10 partial correlations, say, $\rho_1 > \ldots > \rho_{10}$, would be $1/10! \approx 2.76\text{e-}7$. Thus, obtaining such a small probability from a proportion of draws we would at least need 1e10 draws, which is computationally very expensive. For this reason we adopt a more efficient approach.

First, we transform the posterior and prior draws using a Fisher-Z transformation. This transformation changes the space of the partial correlations from (-1,1) to $(-\infty, \infty)$ and makes the joint distribution approximately normal (Mulder, 2016). Note that the order constraints are equivalent for the transformed partial correlations because the Fisher-Z transformation is monotonic. Second we reparameterize the transformed partial correlations to differences, so that an ordering or say, $\rho_1 > \ldots > \rho_5$, can equivalently be written as $(\xi_1, \xi_2, \xi_3, \xi_4) < (0, 0, 0, 0)$, where $\xi_q = \rho_{q+1} - \rho_q$, for $q = 1, \ldots, 4$. As multivariate normality also holds for these difference parameters, we can compute the necessary one-sided probabilities using standard methods, e.g., the 'pmvnorm' function from the R **mvtnorm** package (**?**). In certain specific situations an ordering of parameters cannot be translated to a unique set of differences, e.g., $(\rho_1, \rho_2) > (\rho_3, \rho_4)$. In this case we make use of the R package **bain** (**?**). These methods for computing probabilities for confirmatory testing of order hypotheses in GGMs are implemented in the R package **BGGM**.

Appendix D. Node descriptions

Table D.4: Node descriptions. Clusters: (1) Re-experiencing symptoms, (2) Avoidance, (3) Hyper-arousal.

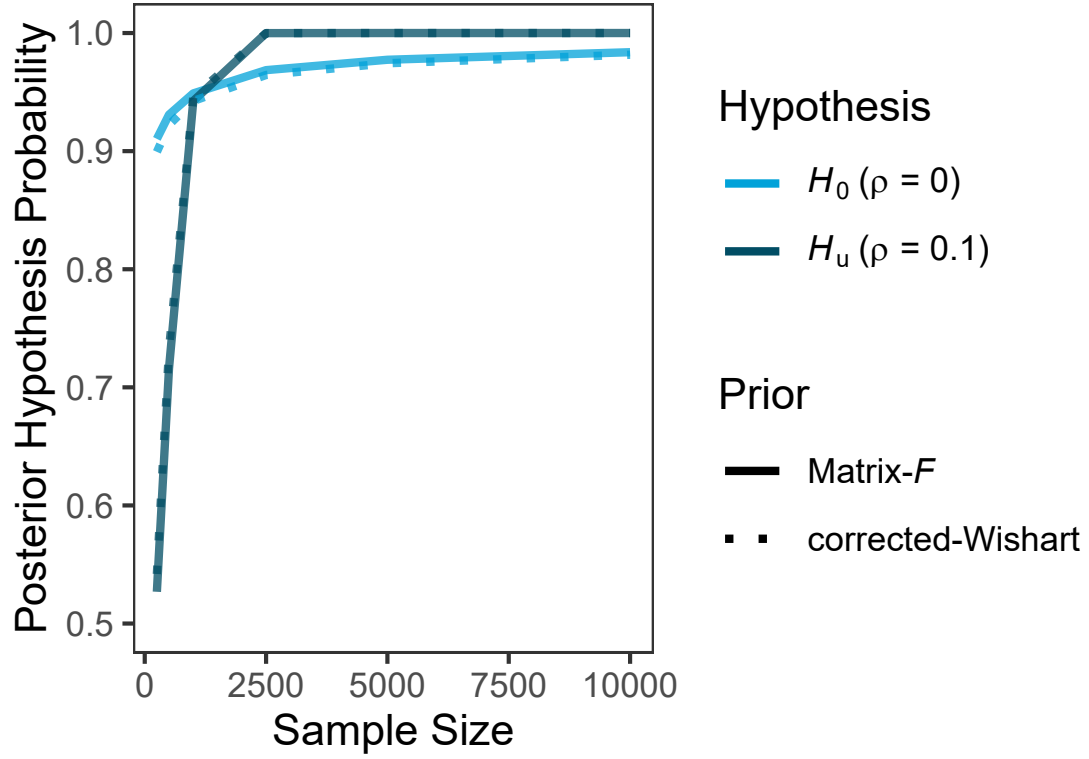| Node	Symptom	Cluster
1	Intrusive thoughts	1 (pink)
2	Nightmares	1 (pink)
3	Avoidance of thoughts	2 (green)
4	Avoidance of situations	2 (green)
5	Sleep problems	3 (blue)
6	Hyper-arousal	3 (blue)

Appendix E. Figures

Figure E.4: Simulation results (Section 3.4.2). The posterior hypothesis probabilities correspond to the true hypotheses $\mathcal{H}_0$ ($\rho = 0$) or $\mathcal{H}_u$ ($\rho \neq 0$). The corrected-Wishart prior approximates the matrix$-F$ prior ($\delta = 2$) with the correction given in (12).