**On unreplicable inferences in psychopathology symptom networks
and the importance of unreliable parameter estimates**

Miriam K. Forbes¹, Aidan G. C. Wright², Kristian E. Markon³, & Robert F. Krueger⁴

¹ Centre for Emotional Health, Department of Psychology, Macquarie University, Sydney, NSW, Australia, 2109

Email: miri.forbes@mq.edu.au

² Department of Psychology, University of Pittsburgh, Pittsburgh, PA, 15260.

Email: aidan@pitt.edu

³ Department of Psychological and Brain Sciences, University of Iowa, Iowa City, IA, 52242.

Email: kristian-markon@uiowa.edu

⁴ Department of Psychology, University of Minnesota, Minneapolis, MN, USA, 55455.

Email: krueg038@umn.edu

Address correspondence to:

Miriam Forbes, PhD
Centre for Emotional Health,
Department of Psychology,
Macquarie University,
Sydney, NSW, Australia, 2109
Email: miri.forbes@mq.edu.au
Phone: +61298509436

Disclosure statement: The authors declare that they have no conflict of interest

Funding details: This work was supported in part by a Macquarie University Research Fellowship; the U.S. National Institutes of Health under grants L30 MH101760 R01AG053217, and U19AG051426; and by the Templeton Foundation.

Draft preprint version 2, November 11th 2020. This paper is not the version of record.

Abstract

We recently wrote an article comparing the conclusions that followed from two different approaches to quantifying the reliability and replicability of psychopathology symptom networks. Two commentaries on the article have raised five core criticisms, which are addressed in this response with supporting evidence. 1) We did not over-generalise about the replicability of symptom networks, but rather focused on interpreting the contradictory conclusions of the two sets of methods we examined. 2) We closely followed established recommendations when estimating and interpreting the networks. 3) We also closely followed the relevant tutorials, and used examples interpreted by experts in the field, to interpret the bootnet and NetworkComparisonTest results. 4) It is possible for statistical control to increase reliability, but that does not appear to be the case here. 5) Distinguishing between statistically significant versus substantive differences makes it clear that the differences between the networks affect the inferences we would make about symptom-level relationships (i.e., the basis of the purported utility of symptom networks). Ultimately, there is an important point of agreement between our article and the commentaries: All of these applied examples of cross-sectional symptom networks are demonstrating unreliable parameter estimates. While the commentaries propose that the resulting differences between networks are not genuine or meaningful because they are not statistically significant, we propose that the unreplicable inferences about the symptom-level relationships of interest fundamentally undermine the utility of the symptom networks.

Keywords: network analysis; network theory of mental disorders; psychopathology symptom networks; replicability crisis

On unreplicable inferences in psychopathology symptom networks and the importance of unreliable parameter estimates

In Forbes, Wright, Markon, and Krueger (2019a), we compared the conclusions that followed from two different approaches to quantifying the reliability and replicability of psychopathology symptom networks: 1) the methods proposed by proponents of symptom networks that are widely used in the literature (i.e., lasso regularisation and the *bootnet* and *NetworkComparisonTest* [NCT] packages in *R*; Epskamp et al., 2018; Epskamp & Fried, 2018), and 2) a set of ‘direct metrics of replicability’ that examined the detailed network characteristics that are interpreted in the literature (e.g., the presence, absence, sign, and strength of individual edges). We compared the conclusions of these approaches based on their application to two networks of depression and generalized anxiety symptoms that were estimated using data collected one week apart in a longitudinal observational study, and to four networks of PTSD symptoms that were the focus of a study on between-sample network replicability led by an expert in the field (Fried et al., 2018). We found that the popular methods tended to obscure evident differences in the networks at the level they are interpreted (e.g., suggesting the networks were accurately estimated, stable, and considerably similar when 39-49% of the estimated edges were unreplicated within each pair of networks).

Here we respond to the five core criticisms raised in two recent commentaries on that article (Fried, van Borkulo, & Epskamp, 2020; Jones, Williams, & McNally, 2020) and highlight what we see as the main point of agreement: The cross-sectional symptom networks are composed of unreliable parameter estimates. Fried et al. (2020) and Jones et al. (2020) propose that the resulting differences between networks are not “genuine” (Jones et al., 2020 p. 3) or “meaningful” (Fried et al., 2020, p. 2) because they are not statistically significant. We conclude our response with discussion on how unreplicable inferences about the symptom-level relationships of interest fundamentally undermine the utility of the symptom networks—particularly in the typical application of these methods wherein a study interprets the results of a single network, making it impossible to know which inferences to take at face value.

Criticism #1: Conclusions about symptom network replicability do not follow from investigations of empirical data; simulations are necessary to understand how methods perform.

Both commentaries rightly argue that it is not appropriate to make broad generalisations about how a statistical method performs (e.g., generalising about the replicability of all symptom networks) on the basis of an isolated application to empirical data. Unfortunately, the commentaries also denounce the target article on these grounds based on a misrepresentation of the aims and conclusions of our article. Specifically, both commentaries proposed that Forbes et al. (2019a) focused on evaluating and criticising the replicability of symptom networks (Fried et al., 2020 pp. 1, 5; Jones et al., 2020, pp. 1, 2). In fact, as described above, the target article focused on comparing the conclusions of two approaches to quantifying the reliability and replicability of symptom networks, discussing the potential reasons they came to contradictory conclusions as well as the implications of the differing conclusions. We were careful not to over-generalise; by contrast Jones et al. (2020) asserted that “extant network analytic methods for assessing replicability [showed] that network characteristics are generally stable and robust” (p. 2) and repeatedly endorsing the conclusion that “Psychopathology networks replicate very well.” (pp. 2, 6), based on isolated applications of the methods to empirical data.

While we were careful not to over-generalise about the replicability of psychopathology symptom networks in Forbes et al. (2019), we did infer that these patterns of contradictory conclusions seem likely to generalise to other applications of lasso regularisation, bootnet, and NCT in psychopathology symptom networks. We also do believe that the patterns of unreplicable network parameters seen here are likely to generalise to other symptom networks estimated in similar data. While thorough and realistic simulations would be informative for understanding whether and under what circumstances we might be able to expect the symptom-level relationships in one network to be replicable, we have now seen substantial instability in network parameters across a variety of empirical applications. For example, we have found network parameters to vary substantially between samples, within samples over time, and between random split-halves and bootstrapped subsamples of a sample (e.g., Forbes et al., 2017a, 2017b, 2019a). These applications in empirical data are illuminating because they consistently converge on parameter instability where it would not be expected on theoretical grounds: In the context of network theory, it is hard to parse

the instability of network parameters, as we are not likely to expect meaningful differences in the symptom-level relationships of interest within a sample over the course of a week, or in random subsamples of the same group of participants, for example. Similarly, it is hard to know how to interpret—in a substantive sense—the variation in network parameters based on the specific details of estimation methods used (e.g., Ising models vs. relative importance networks vs. directed acyclic graphs; regularized vs. unregularized networks; Spearman vs. polychoric correlations; e.g., Epskamp, Kruis, & Marsman, 2017; Forbes et al., 2017a; Fried et al., 2020). Overall, we propose that investigations in empirical data have played an important role in informing our understanding of how this relatively new set of methods perform in data that are typical of our field.

Criticism #2: We did not follow established recommendations when estimating networks, underestimating replicability

Fried et al. (2020) stated that our use of polychoric correlations when estimating the depression and anxiety symptom networks departed from established best practice and resulted in less stable and replicable parameter estimates than if we had used Spearman correlations (pp. 1, 3-4, 5). We were somewhat¹ surprised to read this, as we closely followed the tutorials written by authors of the commentary on estimating and interpreting regularized partial correlation networks (Epskamp & Fried, 2018) and their accuracy (Epskamp et al., 2018) to ensure we followed best practice. Both tutorials specifically recommend using polychoric correlations for networks based on ordinal data (e.g., Epskamp et al., 2018, pp. 198, 202, 209; Epskamp & Fried, 2018, pp. 2, 7), which is also the default in the *qgraph* and *bootnet* packages (Epskamp et al., 2012, 2018).

The three steps recommended in *A Tutorial on Regularized Partial Correlation Networks* (Epskamp & Fried, 2018) are:

- 1) Use polychoric correlations as input. For example: “We estimate regularized partial correlation networks via the Extended Bayesian Information Criterion (EBIC) graphical lasso (Foygel & Drton, 2010), using polychoric correlations as input when data are ordinal. We detail advantages of this methodology, an important one being that it can be used with ordinal variables that are very common in psychological research” (p. 2);

¹ This point mirrors criticisms of our earlier work on symptom network replicability (Forbes et al., 2017a) in which we closely followed the most recent tutorial that was available at the time for estimating symptom networks—even using the same code and data as Borsboom and Cramer (2013). Borsboom et al. (2017) claimed that “inadequacy of the data and analyses” invalidated our results (p. 997), but used those same data and analyses as the basis of their conclusion that “network models replicate very well” (p. 990).

- 2) When using polychoric correlations as input, be wary if the network is “densely connected (i.e., many edges) including many unexpected negative edges and many implausibly high partial correlations (e.g., higher than 0.8)” (p. 13) or “some edges are extremely high and/or unexpectedly negative” (p. 14); and
- 3) If either of the patterns described in 2) occur, check whether networks based on polychoric correlations “look somewhat similar” or “very different”, respectively, to networks based on Spearman correlations (p. 14).

We followed these steps and did not see either (a) a densely connected network, or (b) a pattern of implausibly high or unexpected negative edges, so continued with the recommended default of using polychoric correlations. It seems that the statement in the commentary that we “did not follow established recommendations when estimating networks” (p. 1) contradicts these established recommendations. Naturally, we understand that methods evolve and are refined over time, but patterns of making and then rapidly abandoning methodological recommendations seem to only raise further questions about the robustness of results in the literature.

Criticism #3: We misinterpreted the bootnet and NCT results

Fried et al. (2020) also stated that three sets of conclusions regarding bootnet and NCT results did not follow from our investigation (p. 2). First, that we concluded the interpretation guidelines in the bootnet tutorial (Epskamp et al., 2018) err towards indicating stability and interpretability in the networks, and that bootnet results tended to suggest the networks were accurately estimated. Specifically, Fried et al. stated that this is inaccurate because “*bootnet* results clearly indicate lack of stable estimates in study 1. The CS coefficient [the only *bootnet* output with an objective cut-off provided for interpretation]...was 0.13 for the depression and anxiety networks, implying that the centrality order is unstable and should thus not be interpreted” (p. 2). However, the logic of this argument overlooks that there are four sets of tests provided by bootnet to quantify network accuracy: the CS-coefficient, bootstrapped edge CIs, and bootstrapped difference tests for edge weights and symptom centrality estimates. It was the latter three that we discussed in our paragraph on the bias in the guidelines for interpreting bootnet results (Forbes et al., 2019a, pp. 14-15). Further, the bootnet tutorial emphasises the bootstrapped edge CIs as the single most important method that “should always be performed” when assessing network accuracy (Epskamp et al., 2018, p. 199); it was these results together with the bootstrapped edge difference tests that led us to

conclude that the depression and anxiety network edges were sufficiently stable to interpret (Forbes et al., pp. 8-9), in line with the interpretation guidelines in the two tutorials mentioned above (Epskamp et al., 2018; Epskamp & Fried, 2018). We did not mention the *CS*-coefficient at all in inferring that the networks were accurately estimated. The *CS*-coefficient relates only to interpreting differences in symptom centrality where a result below .25 indicates that any apparent differences should not be interpreted as meaningful (Epskamp et al., 2018, p. 200). This is how we interpreted the value of the *CS*-coefficient in our analyses: “the *CS*-coefficient was below the minimum recommended cut-off at both waves ($CS(0.7) = .13$)... This implies that the apparent differences among symptoms in their standardized strength centrality are not interpretable.” (p. 9).²

Second, Fried et al. (2020) rightly pointed out that it was a misnomer to label edges with bootstrapped CIs that did not span zero “bootnet-accurate”. The reason we examined these edges separately is because they were the only robust edges (i.e., consistently estimated with the same sign) in the bootstrapped networks. All other edges were variably estimated as positive, absent, or negative even within bootstrapped subsamples of the same data. We therefore might expect this robust subset of edges in each network to be the most reliable and replicable between samples, or over time. However, Fried et al. highlighted that these edges are not necessarily “accurate”—in part because they often had large confidence intervals, indicating unreliability (see Maxwell et al., 2008 for a discussion of the role of precision in parameter accuracy). Correspondingly, in the four PTSD networks we examined, 24-38% of these edges in each network were unreplicated (absent altogether) in at least one of the other three networks. Perhaps it would be better to call these edges “bootnet-robust”. Regardless of the label we use, it is the instability of these edges between networks that is the point, and a considerable concern when interpreting a symptom network in a study that relies on a single sample, which is how most of the literature proceeds.

Third, Fried et al. (2020) suggest that the NCT omnibus test may not have had sufficient power to detect differences between the depression and anxiety networks, as the simulations we cited in van Borkulo et al. (under review) were based on continuous data and independent samples whereas the depression and anxiety networks were based on ordinal data and repeated measures of the same sample one week apart. We would typically expect to have greater power to detect group

² In fact, finding strength centrality estimates to be unreliable in the depression and anxiety symptom networks was the only result where bootnet and the direct metrics of replicability converged in their interpretation, which we highlighted on page 11.

differences in repeated measures data, but in order to increase power to detect differences between the networks, we can also examine the NCT individual edge invariance tests before applying the Holm-Bonferroni correction for multiple comparisons. Before correcting for multiple comparisons, four of the edges were deemed significantly different between the two networks. These four edges included only two (5%) of the 42 estimated edges that were unreplicated between the two networks, and (ironically) two edges that were estimated with the same sign in both networks (i.e., replicated). Notably, the edges that reversed in sign and the unreplicated “bootnet-robust” edge were not detected as different between the networks, even with this liberal threshold for statistical significance. It seems clear that the NCT is missing substantive differences between networks that fundamentally affect our inferences regarding the symptom-level relationships in each network.

As for Criticism #2, we were somewhat surprised at the assertion that we misinterpreted the bootnet and NCT results, given we closely followed the tutorials and examples in the literature led by the authors of the commentary. We also included the PTSD networks for this reason: They were estimated and interpreted by a leading expert in the field of symptom networks—the first author of Fried et al. (2020)—which removed any bias we might bring into the subjective interpretation of bootnet results. In that study, Fried et al. (2018) highlighted that the edges of the PTSD networks were accurately estimated based on the bootnet results (p. 342), and repeatedly emphasised the ‘considerable’ similarities between the networks (pp. 335, 343, 344, 346)—noting that only three edges differed substantially across the four samples “whereas others were similar or identical across networks” (p. 344). In contrast to this perspective, we found that 75 (66%) of the edges were inconsistently estimated between the four networks (i.e., variably positive, absent, or negative; see Figure 1) leaving only a third (34%) of the edges consistently estimated with the same sign. In fact, the three edges that were highlighted as “the most pronounced differences” (p. 344) between the networks (1–3, 1–4, and 15–16 in Figure 1) were actually among the minority of edges consistently estimated with the same sign across the four networks, only varying in their estimated weight. This means that the 75 edges that were inconsistently estimated as positive, negative or absent between the four networks were all deemed “similar or identical across networks” by Fried et al. (2018, p. 344). As mentioned above and explored further below, these differences between networks fundamentally affect our inferences regarding the symptom-level relationships in each network, but were obscured when relying on the methods proposed by proponents of symptom networks.

Moreover, it was only through the inclusion of four samples in this study that the extent of the instability in the networks was evident; the majority of symptom network studies include only a single sample (Robinaugh et al., 2020) where it becomes impossible to know which parameters could reasonably be expected to generalise to other samples.

Figure 1. Inconsistently estimated edges among the four PTSD symptom networks. Red edges were inconsistently present/absent, black edges reversed in sign, and dashed edges are negative. Fried et al. (2018) concluded that the four networks were considerably similar, and only three edges (1–3, 1–4, and 15–16) differed substantially “whereas other edges were similar or identical across networks” (p. 344). Ironically, those three edges were among the minority that were consistently estimated with the same sign in all four networks. Overall, two-thirds (66%) of the edges were variably estimated as positive, absent, or negative between the networks.

Criticism #4: Statistical control only reduces reliability when used inappropriately with regard to the underlying causal model

Jones et al. (2020) highlight the important point that statistical control can increase reliability when it is appropriate for the underlying causal model (e.g., controlling for recent physical activity and caffeine consumption when measuring an individual's basal blood pressure, p. 3). We agree that statistical control does not inherently lead to unreliability in parameter estimates. Rather, we argue that indiscriminately including all observed variables of interest as covariates is unlikely to increase the reliability of the resulting highly partialled correlations (e.g., the population value of the partial correlation matrix will change depending on the covariates included; Raveh, 1985). This is especially the case when variables are conceptually part of the same latent constructs with large amounts of shared variance, as is often the case with symptom measures of psychopathology (Lynam et al., 2006).

To see whether statistical control is increasing the reliability of our parameter estimates here, we can examine whether the confidence intervals of the parameter estimates are narrower or wider when comparing the unpartialled (zero-order) correlation for each symptom pair to the fully partialled correlation that underlies the edge weight for those symptoms in the depression and anxiety symptom networks.³ The confidence intervals for the partial correlations were 20-21% wider on average, and 89-92% of the confidence intervals were wider for the partial correlation, compared to the corresponding zero-order correlation at each wave. The partial correlations also tended to be less reliable over time. For example, the average change score was 62% larger between waves for the partial correlations, compared to the zero-order correlations, and 67% of the change scores were larger for the change in the partial correlation between waves compared to change in the zero-order correlation for the same symptom pair. Finally, when we take into account that the partial correlations were smaller in magnitude and examine these change scores as a proportion of the correlations at baseline: partial correlations ($|p| \geq .01$ at wave 1) changed by 115% on average, compared to 8% for zero-order correlations. Overall, it seems likely we are in a scenario more akin to measuring 'lightning seen by people who are deaf' described by Jones et al. (2020, p. 3), rather than homing in on precise estimates of symptom-level relationships.

³ Based on the findings reported by Fried et al. (2020) that the network edges based on Spearman correlations were more reliable and replicable, we estimated 95% confidence intervals based on normal theory using the *psych* package in R (R Core Team, 2019; Revelle, 2019) for zero-order and fully partialled Spearman correlations.

Criticism #5: Statistical tests are necessary for determining whether differences between networks are “genuine” (Jones et al., 2020, p. 5) or “meaningful” (Fried et al., 2020, p. 2) because they take sampling variability into account

A key point in both commentaries is emphasising the importance of taking sampling variability into account in comparing the parameter estimates within and between symptom networks—for example, examining whether parameters’ confidence intervals are overlapping. Specifically, Jones et al. (2020, p. 3) argued:

“[Forbes et al.’s] direct metrics conflate sampling variability and true variability, making them uninterpretable. They regard any difference in the presence or direction of an edge between the two networks as a genuine difference between them. ... We suspected that most of these “changes” between networks arose from ordinary fluctuation between the samples.”

Similarly, Fried et al. (2020) argued:

“In their paper, FWMK analyze and visualize differences in point estimates in detail, and show that statistical tests provided by bootnet and NCT arrive at different conclusions than the authors’ inferences of point estimates. For example, they state: “all NCT results indicated that the depression and anxiety symptom networks had no significant differences when *in fact* they had a multitude of differences” (p. 15, our highlight). This inference is no different than concluding that the *t*-test for neuroticism in Figure 1 reaches a non-significant result when *in fact* point estimates of neuroticism differ across samples (Figure 1)—it ignores sampling variability.” (p. 2, referring to a figure of a box plot depicting two slightly different group means).⁴

In considering these arguments, it is important to distinguish between *statistically significant* versus *substantive* differences between networks. Indeed, this distinction is the basis of the two approaches we compared for quantifying the reliability and replicability of the symptom

⁴ This excerpt from Fried et al. (2020) is also a prototypical strawman argument, based on misrepresenting the original argument. The quote from our paper should read “all NCT results indicated that the depression and anxiety symptom networks had no significant differences when in fact they had a multitude of differences **that fundamentally affected the interpretation of the networks**: Edges varied substantially in weight over time, were often absent altogether in one of the networks, and even occasionally reversed in sign between waves.” (p. 15, emphasis added). The meaning of our argument is clear when the full quote is presented: We are interested in differences that fundamentally affect the interpretation of the networks, not minor differences in point estimates.

networks: bootnet and NCT focus on identifying statistically significant differences between network parameters in bootstrapping and permutation testing frameworks. Both commentaries suggest this type of approach is essential to determine whether apparent differences in network parameters are meaningful. By contrast, the direct metrics of replicability we used focus on identifying substantive differences in the networks *that affect the inferences we would make about symptom-level relationships* (i.e., the basis of the purported utility of symptom networks).

Figure 2. The presence versus absence and sign of an edge affect the inferences we make based on a symptom network. This example shows potential interpretations of an edge between the symptoms of insomnia and fatigue estimated in three networks when the edge is positive, absent, and negative (cf. Forbes et al., 2017a). Even if the differences in this edge between networks are not statistically significant, they lead to different inferences regarding symptom-level dynamics (i.e., the purported utility of symptom networks), and thus represent *substantive* differences.

Consider Figure 2, which depicts the kind of inferences that might be made based on an edge that is variably estimated as present, absent, and negative in three networks (cf. Forbes et al., 2017a). If the networks are estimated using lasso regularisation, the presence and sign of each edge are deemed interpretable and important features of the networks, regardless of the width of the bootstrapped edge CIs (Epskamp et al., 2017; p. 200). In Network 1 the positive edge might thus lead to the inference that difficulty sleeping leads to fatigue; in Network 2 the conditionally dependent symptoms might lead to the inference that insomnia and fatigue share a common cause (e.g., depressed mood); and in Network 3 the negative edge might lead to the inference that being

alert (i.e., - fatigue) makes it harder to fall asleep. Whether or not the differences in this edge are deemed statistically significant between networks, these differences are substantively important because they fundamentally affect our interpretation of the symptom-level relationships. Put simply, the inferences we would draw based on each network are not replicating. This kind of difference is *not* analogous to small fluctuations in an edge of weight 0.005 (no edges in the depression and anxiety networks were this weak), or to slight differences in group means or factor loadings, as proposed in both commentaries (Fried et al., 2020, p. 2; Jones et al., 2020, p. 3). These examples are neither statistically nor substantively different.

Overall, the central message of our paper was that the current “powerful suite of methods to perform tests on the stability and replicability of network analyses” (Jones et al., 2020 p. 6) did not identify important points of instability and non-replication between the networks. A likely explanation for the contradictory conclusions of the two approaches is that the expected variability of the parameter estimates is high. Unreliability in the edge parameters manifests in the sampling variability and uncertainty in parameter estimates highlighted in both commentaries, which in turn leads to large confidence intervals that often overlap and an apparent absence of statistically significant differences between networks. In this context, even networks with a high proportion of unreplicated edges can be deemed ‘considerably similar’ when relying on the popular suite of statistical tests (e.g., Fried et al., 2018).

An important point of agreement, and concluding thoughts

Fundamentally, the target article and commentaries are not disagreeing about whether there are differences between the network edges, but rather on whether these differences are genuine, important, or meaningful. As discussed above, the answer to this question is often clear.

The fact that the unreplicable inferences *between* symptom networks stem from unreliable parameter estimates *within* symptom networks is a particular concern given the large majority of the quantitative literature in this field comprises studies that estimate a single symptom network (Robinaugh et al., 2020). How are we supposed to interpret these networks in isolation? For example, how confident can we be that the paths we find in one network will generalise beyond our specific sample, measure, or set of symptoms? What do we do when the edges in a network vary depending on whether we use Spearman or polychoric correlations? How do we interpret edges that appear, disappear, and reverse in sign when we look at a different sample—or the same sample over

time, or even random subsamples within the same group of people? If these differences are not deemed statistically significantly different, which should we trust when a path does not show up across samples?

The utility of symptom networks rests on generating generalisable inferences about the relationships among symptoms of psychopathology, but we have not yet seen a single set of robust or replicable insights arising from this literature (e.g., Forbes et al., 2019b). It is surprising that prominent rhetoric in the field continues to contend that unreplicable inferences are not important, and continues to promote the reliability, replicability, and general infallibility of symptom networks in the absence of sufficient evidence to support such conclusions (e.g., Jones et al., 2020, Borsboom et al., 2018).

References

- Borsboom, D., & Cramer, A. O. (2013). Network analysis: an integrative approach to the structure of psychopathology. *Annual Review of Clinical Psychology*, 9, 91-121.
<https://doi.org/10.1146/annurev-clinpsy-050212-185608>
- Borsboom, D., Fried, E. I., Epskamp, S., Waldorp, L. J., van Borkulo, C. D., van der Maas, H. L. J., & Cramer, A. O. J. (2017). False alarm? A comprehensive reanalysis of “Evidence that psychopathology symptom networks have limited replicability” by Forbes, Wright, Markon, and Krueger (2017). *Journal of Abnormal Psychology*, 126(7), 989–999.
<https://doi.org/10.1037/abn0000306>
- Borsboom, D., Robinaugh, D. J., The Psychosystems Group, Rhemtulla, M., & Cramer, A. O. (2018). Robustness and replicability of psychopathology networks. *World Psychiatry*, 17(2), 143. <https://dx.doi.org/10.1002%2Fwps.20515>
- Epskamp, S., Borsboom, D., & Fried, E. I. (2018). Estimating psychological networks and their accuracy: A tutorial paper. *Behavior Research Methods*, 50(1), 195-212.
<https://doi.org/10.3758/s13428-017-0862-1>
- Epskamp, S., Cramer, A., Waldorp, L., Schmittmann, V. D., & Borsboom, D. (2012). qgraph: Network visualizations of relationships in psychometric data. *Journal of Statistical Software*, 48 (1), 1-18. <http://dx.doi.org/10.18637/jss.v048.i04>
- Epskamp, S., & Fried, E. I. (2018). A tutorial on regularized partial correlation networks. *Psychological methods*, 23(4), 617. <https://psycnet.apa.org/doi/10.1037/met0000167>
- Epskamp, S., Kruis, J., & Marsman, M. (2017). Estimating psychopathological networks: Be careful what you wish for. *PloS one*, 12(6), e0179891.
<https://doi.org/10.1371/journal.pone.0179891>
- Forbes, M. K., Wright, A. G., Markon, K. E., & Krueger, R. F. (2017a). Evidence that psychopathology symptom networks have limited replicability. *Journal of Abnormal Psychology*, 126(7), 969. <https://doi.org/10.1037/abn0000276>
- Forbes, M. K., Wright, A. G. C., Markon, K. E., & Krueger, R. F. (2017b). Further evidence that psychopathology networks have limited replicability and utility: Response to Borsboom et al. (2017) and Steinley et al. (2017). *Journal of Abnormal Psychology*, 126(7), 1011–1016.
<https://psycnet.apa.org/doi/10.1037/abn0000313>

- Forbes, M. K., Wright, A. G., Markon, K. E., & Krueger, R. F. (2019a). Quantifying the reliability and replicability of psychopathology network characteristics. *Multivariate behavioral research*, 1-19. <https://doi.org/10.1080/00273171.2019.1616526>
- Forbes, M. K., Wright, A. G., Markon, K. E., & Krueger, R. F. (2019b). The network approach to psychopathology: promise versus reality. *World Psychiatry*, 18(3), 272. <https://dx.doi.org/10.1002%2Fwps.20659>
- Fried, E. I., Eidhof, M. B., Palic, S., Costantini, G., Huisman-van Dijk, H. M., Bockting, C. L., ... & Karstoft, K. I. (2018). Replicability and generalizability of posttraumatic stress disorder (PTSD) networks: A cross-cultural multisite study of PTSD symptoms in four trauma patient samples. *Clinical Psychological Science*, 6(3), 335-351. <https://doi.org/10.1177%2F2167702617745092>
- Fried, E. I., van Borkulo, C. D., & Epskamp, S. (2020). On the Importance of Estimating Parameter Uncertainty in Network Psychometrics: A Response to Forbes et al.(2019). *Multivariate Behavioral Research*, 1-6. <https://doi.org/10.1080/00273171.2020.1746903>
- Jones, P. J., Williams, D. R., & McNally, R. J. (2020). Sampling Variability is not Nonreplication: A Bayesian Reanalysis of Forbes, Wright, Markon, & Krueger. <https://doi.org/10.31234/osf.io/egwfi>
- Lynam, D. R., Hoyle, R. H., & Newman, J. P. (2006). The perils of partialling: Cautionary tales from aggression and psychopathy. *Assessment*, 13(3), 328-341. <https://doi.org/10.1177/1073191106290562>
- Maxwell, S. E., Kelley, K., & Rausch, J. R. (2008). Sample size planning for statistical power and accuracy in parameter estimation. *Annual Review of Psychology*, 59, 537-563. <https://doi.org/10.1146/annurev.psych.59.103006.093735>
- R Core Team (2019). R: A language and environment for statistical computing. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Raveh, A. (1985). On the use of the inverse of the correlation matrix in multivariate data analysis. *The American Statistician*, 39(1), 39-42. doi: 10.2307/2683904
- Revelle, W. (2019). *psych: Procedures for Psychological, Psychometric, and Personality Research*. Northwestern University, Evanston, Illinois. R package version 1.8.12, <https://CRAN.R-project.org/package=psych>.

Robinaugh, D. J., Hoekstra, R. H., Toner, E. R., & Borsboom, D. (2020). The network approach to psychopathology: a review of the literature 2008–2018 and an agenda for future research. *Psychological Medicine*, 50(3), 353-366. <https://doi.org/10.1017/S0033291719003404>

van Borkulo, C., Boschloo, L., Kossakowski, J., Tio, P., Schoevers, R., Borsboom, D., & Waldorp, L. (under review). Comparing network structures on three aspects: A permutation test.

Retrieved from https://www.researchgate.net/profile/Claudia_Van_Borkulo/publication/314750838_Comparing_network_structures_on_three_aspects_A_permutation_test/links/58c55ef145851538eb8af8a9/Comparing-network-structures-on-three-aspects-A-permutation-test.pdf