

Bayesian Bootstrapped Correlation Coefficients

Josue E. Rodriguez^a and Donald R. Williams^a

^aUniversity of California, Davis

Abstract ■ We propose the Bayesian bootstrap (BB) as a generic, simple, and accessible method for sampling from the posterior distribution of various correlation coefficients that are commonly used in the social-behavioral sciences. In a series of examples, we demonstrate how the BB can be used to estimate Pearson's, Spearman's, Gaussian rank, Kendall's τ , and polychoric correlations. We also describe an approach based on a region of practical equivalence to evaluate differences and null associations among the estimated correlations. In addition, we have implemented the methodology in the R package `BBcor` (<https://cran.r-project.org/web/packages/BBcor/index.html>). Example code and key advantages of the proposed methods are illustrated in an applied example.

Keywords ■ Bayesian bootstrap, correlation, ordinal, credible interval.

jerrodriguez@ucdavis.edu

10.20982/tqmp.18.1.p039

Acting Editor ■ Denis Cousineau (Université d'Ottawa)

Reviewers

■ Two anonymous reviewers

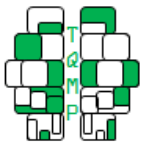
Introduction

Correlation coefficients lie at the heart of research in the social-behavioral sciences (Chen, Smithson, & Popovich, 2002; Cohen, Cohen, West, & Aiken, 2013). They quantify the degree of association between variables, where hypotheses are often posited as correlational statements such as “there is a positive association between IQ and educational attainment.” The most frequently used variant is the Pearson product-moment correlation, or Pearson correlation, that quantifies the strength of the linear association between two variables. Values of 1, -1, and 0, respectively, imply a perfectly positive, perfectly negative, and no relationship.

Although they play a leading role in psychological research, there is surprisingly little work done on estimating common correlation types in a Bayesian framework. To date, the Pearson correlation has received the bulk of attention (e.g., Wetzels & Wagenmakers, 2012; Mulder, 2016; Wagenmakers, Verhagen, & Ly, 2016), but research examining alternative types of correlations is scarce. This is unsurprising because the Pearson correlation is the most frequently used measure of association and it is also trivially estimated, say, by following the separation strategy of Barnard, McCulloch, and Meng (2000) or using the natural conjugate prior for the covariance matrix in a Gaussian model (i.e., the inverse-Wishart). Nevertheless, there are times when researchers would like to estimate a different

type of correlation that may be better suited for their data. For example, Kendall's τ is a popular rank-based correlation method, but was not possible to estimate in a Bayesian framework until only recently (Yuan & Johnson, 2008; van Doorn, Ly, Marsman, & Wagenmakers, 2018). There are a variety of reasons for why this is the case, for instance, due to the lack of an explicit likelihood function and sensible choices for prior distributions (Yuan & Johnson, 2008). Furthermore, polychoric correlations, commonly used for ordinal data, can be challenging to implement and computationally expensive to estimate (e.g., Lawrence, Bingham, Liu, & Nair, 2008). One such approach is the multivariate probit model (e.g., Albert, 1992; Chib & Greenberg, 1998), but this requires sampling latent (Gaussian) data and thresholds, both of which are not straightforward. These methodological challenges have resulted in a lack of software for estimating Bayesian correlations.

To overcome these hurdles, we propose the Bayesian bootstrap (BB, Rubin, 1981) as a simple and flexible approach to obtain a posterior distribution for a correlation matrix. This method is attractive in the sense that it avoids the direct specification of a prior and is straightforward to implement because it is operationally equivalent to the classical bootstrap (Efron, 1979). The key difference between them is that the BB attaches weights to the observed values from a uniform Dirichlet distribution, as opposed to the classical bootstrap that resamples the data. The main benefit of this weighting scheme is that the resulting sam-



ples can be used to approximate the posterior distribution of interest under a noninformative prior (Lo, 1987, 1988; Weng, 1989; Lyddon, Holmes, & Walker, 2019). The motivation behind the BB is nicely summarized in Kim and Lee (2003),

“To circumvent such complications of the full Bayesian analysis, we propose Bayesian bootstrap (BB) procedures which, we believe, are easily accessible to practitioners and at the same time are reliable inference procedures...the BB procedures are conceptually parametric and conceptually simple but retain the flexibility of nonparametric models. Another advantage of the BB procedures is that it is unnecessary to elicit prior information...” (p. 1905)

Because the BB is flexible and does not require a prior to be explicitly specified by the analyst, it can be used to seamlessly estimate virtually any correlation matrix, including Kendall's τ and polychoric correlations. However, the BB remains relatively unknown in psychological contexts despite its simple form and utility with respect to simulating samples from the posterior distribution.

Naturally, a key attraction of the BB is that it shares important properties with traditional Bayesian inference. The benefits of adopting Bayesian approaches have been written about extensively in the psychological sciences (see e.g., Vandekerckhove, Rouder, & Kruschke, 2018, and other articles in that special issue). For instance, analysts commonly want to make statements about which parameter values are the most likely, conditional on the observed data (Kruschke, Aguinis, & Joo, 2012; Kruschke, 2018), but this privilege is reserved for Bayesian methods as opposed to classical inferential techniques. Consequently, adopting a Bayesian approach necessarily results in a posterior distribution, and thus, statements can be made about the probability of specific parameter values, or a range of them (Wagenmakers, Morey, & Lee, 2016; Wagenmakers et al., 2018). Moreover, Bayesian inference allows for quantifying evidence in favor of a null hypothesis as opposed to more classical methods which typically only allow for rejecting or failing to reject the null hypothesis.

Because the Bayesian bootstrap provides a valid posterior, it can be further employed to compare correlations. The problem of comparing correlations from the same sample has received ample attention in the literature (Dunn & Clark, 1969; Steiger, 1980; Meng, Rosenthal, & Rubin, 1992; Raghunathan, Rosenthal, & Rubin, 1996; Zou, 2007; Mulder, 2016), and there are three main cases where comparing correlations is of interest (Krishnamoor-

thy & Xia, 2007): (1) overlapping dependent correlations, (2) non-overlapping dependent correlations, and (3) independent correlations from independent samples. Because the dependence structure is encoded in the posterior distribution, the BB can be employed in all of these situations.

Major Contributions

This work includes three major contributions. First, the Bayesian bootstrap is introduced as a method for approximating posterior distributions for several correlation coefficients. Namely, we describe the Bayesian bootstrap for the Pearson correlation, wherein the Spearman's and Gaussian rank correlations naturally arise as special cases. We further provide formulations to obtain Kendall's and polychoric correlation coefficients. We emphasize that these latter two coefficients, unlike the Spearman's and Gaussian rank correlations, cannot be trivially estimated in a Bayesian framework. Second, an approach is discussed for comparing two or more correlations, possibly with the region of practical equivalence (ROPE) of Kruschke (2018). This allows researchers to go beyond merely estimating correlations to making meaningful comparisons among them (e.g., establishing null associations). Third, to increase the availability of the proposed approach, Bayesian bootstrapped correlations have been implemented in the R package `BBcor`. For users who are unfamiliar with R, we have implemented a Shiny app¹ (Chang et al., 2021). The totality of these contributions places the Bayesian bootstrap into the toolbox of researcher psychologists.

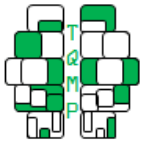
Overview

The outline of this article is as follows. We begin by delineating the Bayesian bootstrap procedure for different correlation types. Here it is shown how estimating correlations with the BB essentially amounts to calculating weighted correlations. Next, we demonstrate how two or more correlations can be compared with the resulting posterior distribution. We then move on to empirical illustrations of the method using two psychological datasets. These examples illustrate the utility of the proposed method in applied settings. We conclude with a brief discussion on the Bayesian bootstrap.

The Bayesian Bootstrap

There are at least three ways to view the Bayesian bootstrap (Kim & Lee, 2003): 1) as an extension of the classical bootstrap, 2) the limit of the full Bayesian posterior as the prior becomes completely uninformative (Gasparini, 1995, Theorem 2), and 3) a distribution that is proportional to the product of the empirical likelihood and an uninforma-

¹The Shiny app can be accessed at tinyurl.com/2nw33cu8



tive prior (Rubin, 1981; Owen, 1990; Lazar, 2003; Choudhuri, 1998). Because in psychology, most analysts are likely to have at least some familiarity with the classical bootstrap, we briefly describe this perspective here. Suppose $Y = (y_1, \dots, y_n)$ is a random sample from an unknown distribution F and we are interested in estimating a functional of F , $T(F)$, say, the expected value of Y . The classical bootstrap entails resampling the data with replacement to obtain Y_1^*, \dots, Y_B^* where B is the number of bootstrap samples. Inferences are then drawn on the basis of $T(F_i^*)$, where F_i^* is the empirical distribution of the i th resampled dataset. Notice that the empirical distribution can be expressed as $F_i^* = \sum_j^n w_j \delta_{Y_j}$ where $n(w_1, \dots, w_n) \sim \text{Multinomial}(n, 1/n, \dots, 1/n)$. The weights w are discrete, considered to be known, and denote the proportion a distinct value of the original data, δ_{Y_j} , arises in the bootstrap sample. By instead considering the weights for each sample to be unknown, continuous, and distributed as $\text{Dirichlet}(1, \dots, 1)$, the resulting empirical distribution F_i^* takes on a smoother shape (see Figure 1 in Rubin, 1981). Technical details of the connection between the Bayesian bootstrap and the usual posterior distribution are given in the appendix. For comprehensive mathematical treatments of the BB, we refer readers to Lo (1987, 1988), Newton and Raftery (1994), and the references therein.

Illustration

To illustrate the process of obtaining a BB posterior, suppose that we have n observations of a random variable Y . The BB generates a posterior probability for each observation y_1, \dots, y_n , where unobserved values have zero posterior probability. Specifically, one BB sample is obtained by drawing n weights from a uniform Dirichlet distribution and attaching them to the data. The generated weights can be interpreted as the probabilities that $Y = y_i$ in each sample (Rubin, 1981). In practice, these weights are easily generated using draws from an exponential distribution (see e.g., Devroye, 1986, p. 594). If this process is repeated S times, then the distribution of all S samples is the BB distribution of Y . More often, however, we are interested in estimating the parameter of a distribution, say, the mean. For each s sample ($s = 1, \dots, S$), the steps for estimating the mean of Y are as follows:

1. Draw n exponential variates

$$z_i^{(s)} \sim \text{Exp}(1), \quad i = 1, \dots, n \quad (1)$$

2. Generate the weights

$$w_i^{(s)} = \frac{z_i^{(s)}}{\sum_{i=1}^n z_i^{(s)}} \quad (2)$$

3. Calculate the weighted sample mean

$$\bar{y}^{(s)} = \sum_{i=1}^n w_i^{(s)} y_i \quad (3)$$

The empirical distribution of $\{\bar{y}^{(1)}, \dots, \bar{y}^{(S)}\}$ is the BB approximation to the posterior of the mean of Y . A visual comparison between a BB distribution and an analytical posterior for this scenario is shown in Figure 1. Note steps 1 and 2 can be merged if a uniform Dirichlet distribution random number generator is directly available, as is the case in many programming platforms. Further, a subscript can be added $\bar{y}_g^{(s)}$ ($g = 1, \dots, G$) in each step to distinguish means between groups. This opens up the possibility to obtain a posterior distribution for mean differences (e.g., $\delta^{(s)} = \bar{y}_1^{(s)} - \bar{y}_2^{(s)}$). In what follows, we demonstrate how these ideas can be harnessed to estimate and compare a variety of correlation coefficients.

Pearson, Spearman's, and Gaussian Rank Correlation Coefficients

Background

The most popular correlation is the Pearson product-moment correlation coefficient, or Pearson correlation, which captures the linear relationship between two variables. When the data are ordinal, it is common to use the nonparametric Spearman's correlation, which is defined as the Pearson correlation between the ranks of two variables and describes their monotonic relationship. Although conceptually easy to understand and compute, using Spearman's correlation results in a small loss of statistical efficiency. A recently proposed alternative is the Gaussian rank correlation (Boudt, Cornelissen, & Croux, 2012). The Gaussian rank correlation is defined as the Pearson correlation of the normalized ranks (i.e., their Van der Waerden scores). The advantage of normalizing the variables prior to computing their correlation is that there is a small gain in statistical efficiency (for Gaussian data) when estimating the monotonic relationship between them. Interestingly, the main difference between the Pearson's, Spearman's, and Gaussian rank correlations is whether the raw, ranked, or normalized rank observations are being correlated. Hence, only a formulation for the Pearson correlation is needed to obtain any of the three correlation types. Note that Rubin (1981) described the BB for a sec:comparingsingle Pearson's correlation, but did not consider the full correlation matrix or other correlation types.

In a Bayesian framework, the Pearson correlation matrix is traditionally estimated by modeling the covariance matrix Σ . To this end, the legacy inverse-Wishart prior has been the de facto standard. Due to its conjugacy, computation can be relatively efficient and thus it is widely imple-

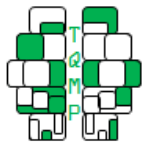
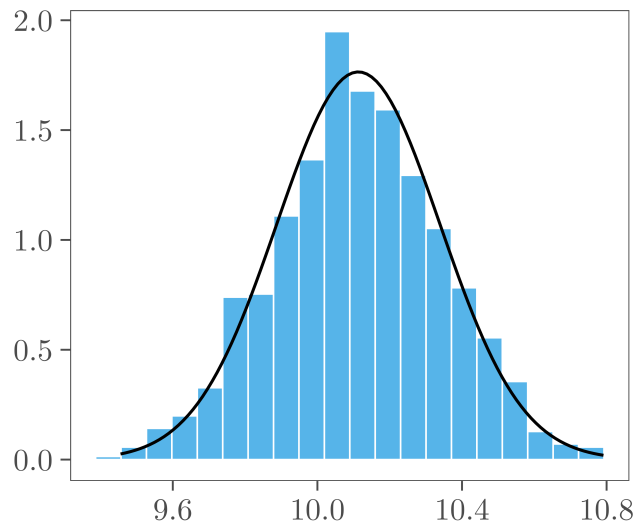


Figure 1 ■ Density histogram of 1000 Bayesian bootstrapped means using the steps outlined in (1) – (3). The black line is the posterior density of the mean resulting from a normal prior with mean equal to zero and standard deviation equal to 10. The sample data ($n = 500$) was generated from a Normal distribution with mean equal to 10 and standard deviation equal to 5.



mented in Bayesian software (e.g., Plummer, 2003). However, the inverse-Wishart prior has been criticized for several reasons: the uncertainty for all variances is controlled by a single degrees of freedom parameter (Barnard et al., 2000), the marginal distribution for the variances have low density near zero (Gelman, 2006), and there is an *a priori* dependence between the resulting correlations and variances (Tokuda, Goodrich, Van Mechelen, Gelman, & Tuerlinckx, 2011). Separation strategies exist to deal with the dependence between the variances and correlations (e.g., Barnard et al., 2000), but suffer from similar problems as the inverse-Wishart. Alternative distributions exist that circumvent these issues, such as the LKJ (Lewandowski, Kurowicka, & Joe, 2009) or matrix- F (Mulder & Pericchi, 2018) prior distributions. Although they are more flexible than the inverse-Wishart, the incurred expense is that they are more computationally complex and, additionally, are not yet widely available in Bayesian software. For instance, the LKJ prior is mostly restricted to programs that interface with Stan (Carpenter et al., 2017) and do not readily provide the full correlation matrix. The matrix- F prior has been implemented for a full correlation matrix, but first requires estimating the partial correlations and thus the prior cannot be placed directly over the correlation matrix (Williams & Mulder, 2019). By instead employing the Bayesian bootstrap, an approximate posterior for the full correlation matrix can be obtained painlessly.

Bayesian Bootstrap Steps

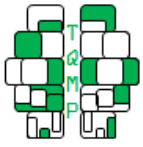
We now describe the necessary ingredients for obtaining Bayesian bootstrapped samples of Pearson, Spearman's, and Gaussian rank correlations. Without a loss of generality, assume \mathbf{Y} to be a mean-centered $n \times p$ data matrix with sample covariance matrix \mathbf{S} . The Pearson correlation matrix for \mathbf{Y} is given by

$$\mathbf{R} = \mathbf{D}^{-\frac{1}{2}} \mathbf{S} \mathbf{D}^{-\frac{1}{2}} \quad (4)$$

$$\mathbf{S} = (n - 1)^{-1} (\mathbf{Y}' \mathbf{Y})$$

where $\mathbf{D}^{-\frac{1}{2}}$ is a diagonal matrix containing the inverse square roots of the diagonal elements of \mathbf{S} and each r_{ij} element of \mathbf{R} indicates the correlation between the i th and j th column of \mathbf{Y} . The Spearman's correlation matrix is obtained when each (i, j) th element in \mathbf{Y} is replaced with its rank, $R(Y_{ij})$. Similarly, if the elements are replaced with their Van der Waerden scores, $\Phi^{-1} \left(\frac{R(Y_{ij})}{n+1} \right)$, where Φ^{-1} denotes the quantile function for a standard normal curve, then \mathbf{R} contains the Gaussian rank correlations.

A simple modification of (4) yields a posterior sample of \mathbf{R} . Mainly, for each s sample, the values drawn from the uniform Dirichlet distribution are used to center the columns of \mathbf{Y} at their *weighted* mean, and then further to obtain a *weighted* covariance matrix \mathbf{S}_w . With this modifi-



cation, a Bayesian bootstrapped sample for \mathbf{R} is obtained by computing

$$\mathbf{R}_w^{(s)} = \left(\mathbf{D}_w^{(s)}\right)^{-\frac{1}{2}} \mathbf{S}_w^{(s)} \left(\mathbf{D}_w^{(s)}\right)^{-\frac{1}{2}} \quad (5)$$

$$\mathbf{S}_w^{(s)} = \left[1 - \sum_{i=1}^n \left(w_i^{(s)}\right)^2\right]^{-1} \left(\mathbf{Y}_w^{(s)'} \mathbf{Y}_w^{(s)}\right) \quad (6)$$

$$\mathbf{Y}_w^{(s)} = \mathbf{Y} \circ \mathbf{w}_*^{(s)} \mathbf{1}_p' \quad (7)$$

where $\mathbf{R}_w^{(s)}$ is a weighted correlation matrix, $\left(\mathbf{D}_w^{(s)}\right)^{-1/2}$ is a diagonal matrix containing the inverse square roots of the diagonal elements of $\mathbf{S}_w^{(s)}$, and $\mathbf{Y}_w^{(s)}$ is a weighted version of the data matrix. The symbol “ \circ ” denotes the Hadamard product, $\mathbf{w}_*^{(s)}$ is an n -dimensional vector with elements $w_{*,i}^{(s)} = \sqrt{w_i^{(s)}}$, and $\mathbf{1}_p$ is a p -dimensional vector containing 1's. If $\mathbf{R}_w^{(s)}$ is computed S times, then the distribution of $\{\mathbf{R}_w^{(1)}, \dots, \mathbf{R}_w^{(S)}\}$ is the BB distribution of \mathbf{R} . Similarly, the BB distribution of each r_{ij} is the empirical distribution of $\{r_{w,ij}^{(1)}, \dots, r_{w,ij}^{(S)}\}$. Notice that computing a posterior sample with the BB requires only a few steps and does not involve explicitly invoking a prior distribution. In this way, the Bayesian bootstrap provides a seamless method for obtaining posterior distributions for the Pearson, Spearman's, and Gaussian rank correlation matrices.

Kendall's Rank Correlation Coefficient

Background

A similar approach can be taken to obtain posterior samples for Kendall's rank correlation coefficient (Kendall, 1938), or Kendall's τ , a widely used measure of association in nonparametric statistics. Like Spearman's correlation, it is a robust measure that captures monotonic relationships between two variables, but has some advantages. It is asymptotically more efficient and has an appealing interpretation. Kendall's τ can be interpreted as follows. Suppose we have n observations for two random variables X and Y . A pair of differences $(x_i - x_j)$ and $(y_i - y_j)$ is said to be concordant if they share the same sign and discordant if they do not. Kendall's τ is obtained by taking the difference between concordant and discordant pairs and dividing this quantity by the number of all possible pairs. When $\tau = 1$ (-1) all pairs of observations are concordant (discordant).

Despite its popularity, there is a dearth of literature on Bayesian inference for Kendall's rank correlation. The main reason for this is that nonparametric tests in Bayesian settings have historically been limited by a lack

of prior distributions and an explicit likelihood function (Yuan & Johnson, 2008) — without which a model cannot be formulated in a Bayesian framework. Recently, van Doorn et al. (2018) developed a method for deriving a posterior distribution for Kendall's τ based on its standardized test statistic T^* . However, this method only considers a single correlation at a time. That is, the full correlation matrix is not readily estimated, which, in turn, prevents easily comparing correlations. In contrast, a Bayesian bootstrap approach to estimating Kendall's τ circumvents this concern because it readily estimate the full correlation matrix.

Bayesian Bootstrap Steps

For the case of X and Y , Kendall's τ is defined as

$$\tau = \frac{\sum_{1 \leq i < j \leq n} \text{sgn}(x_i - x_j) \text{sgn}(y_i - y_j)}{k}, \quad (8)$$

where $k = \frac{n(n-1)}{2}$ is the number of distinct pairs. The above is commonly referred to as τ_A and does not account for ties. When ties are present, the denominator is adjusted to correct for this and is defined as $\sqrt{(k - t_x)(k - t_y)}$ where t_x and t_y denote the number of ties in X and Y , respectively. This version is commonly known as τ_B and because this is the version we consider here, we simply refer to it as τ .

A Bayesian bootstrapped sample for Kendall's rank correlation between X and Y can be computed by first drawing values from a uniform Dirichlet distribution and weighting the numerator to obtain

$$\tau_w^{(s)} = \sum_{1 \leq i < j \leq n} w_i^{(s)} w_j^{(s)} \text{sgn}(x_i - x_j) \text{sgn}(y_i - y_j). \quad (9)$$

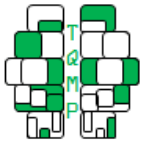
This expression is nice because the connection to the original formulation is clear, but it can be generalized to obtain the full correlation matrix (Pozzi, Di Matteo, & Aste, 2012, pp. 15 – 17). Let \mathbf{Z} be a $k \times p$ matrix where each (l, v) element is associated to $\text{sgn}(y_i^v - y_j^v)$ where $l = 1, \dots, k$, $v = 1, \dots, p$, $i = 2, \dots, n$, and $j = 1, \dots, n - 1$, or in words, a matrix where each element indicates the sign for the difference of the observation pair (i, j) on variable v . With this definition of \mathbf{Z} , a Bayesian bootstrap sample for the Kendall's τ correlation matrix can be obtained as follows

$$\tau_w^{(s)} = \left(\mathbf{D}_w^{(s)}\right)^{-1/2} \mathbf{S}_w^{(s)} \left(\mathbf{D}_w^{(s)}\right)^{-1/2} \quad (10)$$

$$\mathbf{S}_w^{(s)} = \mathbf{Z}^{(s)'} \mathbf{Z}_w^{(s)} \quad (11)$$

$$\mathbf{Z}_w^{(s)} = \mathbf{Z} \circ \mathbf{w}_*^{(s)} \mathbf{1}_p' \quad (12)$$

where $\tau_w^{(s)}$ is the weighted correlation matrix, $\left(\mathbf{D}_w^{(s)}\right)^{-1/2}$ is a diagonal matrix containing the inverse square roots of



the diagonal elements of $\mathbf{S}_w^{(s)}$, and $\mathbf{Z}_w^{(s)}$ is a weighted version of \mathbf{Z} . The k -dimensional vector $\mathbf{w}_*^{(s)}$ contains the elements $w_{*,l}^{(s)} = \sqrt{w_i^{(s)} w_j^{(s)}}$, and $\mathbf{1}_p$ is p -dimensional vector containing 1's. When written this way, Kendall's rank correlation can be conceptualized as a Pearson correlation computed with \mathbf{Z} . If $\tau_w^{(s)}$ is computed S times, then the distribution of $\{\tau_w^{(1)}, \dots, \tau_w^{(S)}\}$ is the BB distribution of τ . Like each $r_{w,ij}$, the BB distribution of each $\tau_{w,ij}$ is their empirical distribution over all S samples.

Polychoric Correlation Coefficient

Background

An important measure of association in the field of psychometrics is the polychoric correlation coefficient (Olsson, 1979; Jöreskog, 1994). Like correlations that describe monotonic relationships, the polychoric correlation is often used with ordinal data. The key difference here is that the ordinal data are considered to be the result of discretizing continuous variables. Accordingly, the polychoric correlation captures the linear association between two latent continuous variables underlying the observed ordinal data. Note that we assume that the joint distribution of the two latent variables is Gaussian, but other distributions can be used (e.g., bivariate t , Albert, 1992).

Getting Bayesian estimates of polychoric correlations can be difficult due to their computational expense. Computing the likelihood requires iteratively sampling from truncated Gaussian distributions and the covariance matrix is typically restricted to be a correlation matrix for identifiability reasons (Albert, 1992; Chib & Greenberg, 1998). Further, nuisance parameters, termed thresholds, must be estimated for each variable. More efficient MCMC algorithms have been developed, for example, by using parameter expansion for data augmentation (Talhouk, Doucet, & Murphy, 2012; Lawrence et al., 2008) or parameterising the precision matrix of the latent variables in terms of the Cholesky decomposition (Webb & Forster, 2008), but these techniques introduce computational complexities of their own and remain unavailable in statistical software (to our knowledge). Thus, for polychoric correlations, the Bayesian bootstrap again provides a relatively simple solution.

Bayesian Bootstrap Steps

For ease of exposition, we focus on estimating the polychoric correlation between two variables, but the following can be applied for the entire correlation matrix. Suppose that two ordinal variables X and Y are expressed in a two-way contingency table with R rows and C columns. That is, there are R levels in X and C levels in Y . If the data is collected on n individuals and classified with re-

spect to the rows and columns, then the cell counts, n_{rc} ($r = 1, \dots, R$, $c = 1, \dots, C$) have respective probabilities π_{rc} . The typical estimation approach is then to assume that the ordinal variables correspond to continuous Gaussian variables ξ and η . The n pairs (ξ_i, η_i) can likewise be placed in an $R \times C$ contingency table using row thresholds $-\infty = a_0 < a_1 < \dots < a_{R-1} < a_R = \infty$ and column thresholds $-\infty = b_0 < b_1 < \dots < b_{C-1} < b_C = \infty$. The relationship between X and ξ is

$$x_i = \begin{cases} 1 & \text{if } \xi_i < a_1 \\ 2 & \text{if } a_1 \leq \xi_i < a_2 \\ \vdots & \\ R & \text{if } a_{R-1} \leq \xi_i \end{cases}, \quad (13)$$

and similarly for Y and η .

The polychoric correlation can then be estimated in two steps (Olsson, 1979). The thresholds are first estimated as

$$a_r = \Phi^{-1} \left(\frac{\sum_{i=1}^n \mathbb{I}(x_i \leq r)}{n} \right), \quad r = 1, \dots, R-1 \quad (14)$$

$$b_c = \Phi^{-1} \left(\frac{\sum_{i=1}^n \mathbb{I}(y_i \leq c)}{n} \right), \quad c = 1, \dots, C-1, \quad (15)$$

where Φ denotes the bivariate standard normal cumulative density function with correlation ρ and the symbol $\mathbb{I}(\cdot)$ denotes the indicator function. Then, the likelihood of the sample

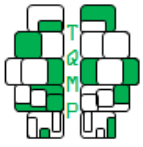
$$\sum_{r=1}^R \sum_{c=1}^C n_{rc} \ln \pi_{rc} \quad (16)$$

is maximized with respect to ρ . Above, n_{rc} is the number of observations in the (r, c) th cell of the contingency table and π_{rc} is the probability that (ξ_i, η_i) belongs to that cell

$$\pi_{rc} = \Phi(a_r, b_c) - \Phi(a_{r-1}, b_c) - \Phi(a_r, b_{c-1}) + \Phi(a_{r-1}, b_{c-1}). \quad (17)$$

The value of ρ that maximizes the log-likelihood is the estimate for the polychoric correlation between X and Y .

A Bayesian bootstrapped sample of the polychoric coefficient can be obtained through a reweighting scheme applied to the $R \times C$ contingency table. To obtain the weighted cell probabilities, the thresholds are first estimated based on the simulated Dirichlet weights (Bailey, Emad, Zhang, Xie, & Sikali, 2018)



$$a_{w,r}^{(s)} = \Phi^{-1} \left(\sum_{i=1}^n w_i^{(s)} \mathbb{I}(x_i \leq r) \right) \quad (18)$$

$$b_{w,c}^{(s)} = \Phi^{-1} \left(\sum_{i=1}^n w_i^{(s)} \mathbb{I}(y_i \leq c) \right). \quad (19)$$

Similarly, the term n_{rc} in (16) is replaced with

$$n_{w,rc}^{(s)} = \sum_{i=1}^n w_i^{(s)} \mathbb{I}(x_i = r) \mathbb{I}(y_i = c). \quad (20)$$

The weighted probabilities for each sample $\pi_{w,rc}^{(s)}$ are computed using the expression in (17), but with the weighted thresholds so that the log-likelihood for each sample is given by

$$\sum_{i=1}^s \sum_{j=1}^r n_{w,rc}^{(s)} \ln \pi_{w,rc}^{(s)}. \quad (21)$$

Finally, the Bayesian bootstrapped sample for the polychoric correlation, $\rho^{(s)}$, is the one that maximizes (21). If this procedure is carried out S times, then $\{\rho^{(1)}, \dots, \rho^{(S)}\}$ is the BB distribution of the polychoric correlation between X and Y .

Comparing Correlations

Once a set of correlations has been estimated, a common next step is to make comparisons among them, say, to determine which association is the largest. This can be done by computing the posterior distribution for comparisons of interest. The main advantage of doing so is that standard deviations (analogous to standard errors) are available in situations where they would otherwise be difficult to obtain (e.g., the difference between two polychoric correlations with the same matrix). Fortunately, the Bayesian bootstrapped posterior distribution can be used to make such comparisons.

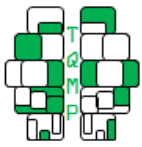
Using the Bayesian bootstrap, the posterior can be obtained for linear combinations of correlations by manipulating the posterior samples of the individual correlations. Say we have estimated a $p \times p$ correlation matrix and are interested in their pairwise differences. Let $\rho^{(s)}$ be a vector containing the s th sample for the $G = p(p-1)/2$ distinct correlations and \mathbf{C} be a matrix of coefficients capturing the pairwise differences. Each element of \mathbf{C} is either a 1, -1 , or 0. A posterior sample for these differences can be obtained by expressing them as a linear combination

$$\delta^{(s)} = \mathbf{C} \rho^{(s)} \quad (22)$$

$$\rho^{(s)} = \begin{bmatrix} \rho_1^{(s)} \\ \rho_2^{(s)} \\ \vdots \\ \rho_G^{(s)} \end{bmatrix}, \quad \mathbf{C} = \begin{bmatrix} 1 & -1 & 0 & \dots & 0 \\ 1 & 0 & -1 & \dots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & 1 & -1 \end{bmatrix}. \quad (23)$$

The distribution of all $\{\delta^{(1)}, \dots, \delta^{(S)}\}$ approximates the posterior distribution for the comparisons between the correlations in ρ . Now, means, standard deviations, and credible intervals can be computed directly for the posterior of δ . The subscripts of the G correlations can denote distinct correlations within the same group, the same correlation for distinct groups, or distinct correlations from distinct groups. Although we focused on pairwise differences here, this idea can be extended to more general linear combinations.

An additional advantage of Bayesian analysis is the ability to “accept” parameter values that provide support either for or against a null hypothesis. For instance, if one wants to conclude that there is no difference between the magnitude of two correlations, then this can be done using a formal procedure such as the region of practical equivalence (ROPE) approach (Kruschke, 2018). The ROPE approach is similar in spirit to a frequentist approach wherein a prespecified parameter value is rejected if it is not covered by a $100(1 - \alpha)\%$ confidence interval. The difference is that a range of parameter values (i.e., a ROPE) is stipulated where values in this range are treated as equivalent to a null value (e.g., 0). Once this region is established and the posterior distribution of δ has been computed, a $100(1 - \alpha)\%$ credible interval (CrI) can be constructed for each comparison. If the computed interval lies entirely inside the ROPE bounds, then the estimated parameter value is treated as equivalent to the null value and conversely, if the interval completely excludes the ROPE, then the null value is rejected. This is because a $100(1 - \alpha)\%$ CrI contains the $100(1 - \alpha)\%$ most probable values (assuming a symmetric distribution). Thus, if the CrI is entirely inside of the ROPE, then we can interpret the parameter value as being practically equivalent to the null value and vice versa. A decision is withheld if there is overlap between the interval and the ROPE. Adopting the framework described above permits researchers to utilize the BB to make meaningful comparisons between associations using a variety of different correlation coefficients with the goals of either parameter estimation, making decisions about a parameter value, or both.



Summary

In this section, we described how posterior distributions for several different correlations can be obtained in a straightforward manner via the Bayesian bootstrap. The central theme was that simulating posterior samples for correlations boils down to repeatedly calculating weighted correlations where the weights are uniform Dirichlet distributed. In each iteration of the bootstrap, the resulting weighted correlation constitutes a draw from the correlation's posterior distribution, and when done repeatedly, the distribution of the calculated statistics approximates the posterior of interest. The main advantage of this method is that posterior inference for correlations can be done “painlessly”. That is, obtaining BB estimates for the correlations does not require specifying a prior distribution or complex sampling techniques. Altogether, the BB provides a powerful tool for approximate Bayesian inference of popular correlation types in social-behavior sciences.

Empirical Application

Below we discuss an empirical example where we illustrate how the `BBcor`² package can be applied to obtain and compare Bayesian bootstrapped correlations in practice. We utilize data that were first analyzed in Šrol, Cavojava, and Mikušková (2021) to compare dependent correlations from the same sample. The data were collected to study the negative social consequences of Covid-19 related conspiracy beliefs. Slovakian participants ($N = 501$) completed survey items measuring their prejudiced and discriminatory views against three social outgroups associated with the pandemic in Slovakia. Specifically, data were collected on negative feelings, social distance, and discriminatory views towards Chinese, Roma, and Italian people. Further, measurements were taken on the degree of belief in general Covid-19 conspiracies (e.g., “Covid-19 is a biological weapon intended to eliminate the overcrowded human population”) and Chinese-specific Covid-19 conspiracies (e.g., “the Chinese created [SARS-CoV-2] as a biological weapon which then got out of hand”). As part of the analysis in this study, the three measures of prejudice and discrimination were each correlated with the measures of conspiracy belief, yielding six correlations per subgroup. The resulting correlations were then compared using Steiger's z -test (Steiger, 1980). For example, the correlation between negative feelings towards Italians and general Covid-19 conspiracy beliefs was compared to the correlation between negative feeling towards Italians and Chinese-specific Covid-19 conspiracies.

There are two details to note here. First, the z -test used

to compare correlations makes the assumption that the underlying data are Gaussian. Second, failing to reject the null hypothesis does not provide support in favor of no difference (i.e., absence of evidence is not evidence of absence). Thus, it may be desirable to use a method of comparison that accommodates a measure of association more appropriate for Likert-type data, such as the data collected (e.g., Kendall's τ), and that allows for statements in favor of the null hypothesis. This can easily be accomplished with the Bayesian bootstrap methodology outlined in this article.

Calculating the correlations

We assume the reader to have some familiarity with the R programming language (R Core Team, 2021). To begin, the `BBcor` package must be installed and loaded, and the data must be read into R.

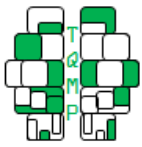
```
# install and load BBcor
install.packages("BBcor")
library(BBcor)
# read in data set
data("srol2021")
str(srol2021)
> 'data.frame': 501 obs. of 11 variables:
> $ neg_feelings_china : int 100 96 75 50...
> $ social_distance_china : num 7 7 5.33...
> $ discrimination_china : int 7 7 5 3 1 6...
> $ neg_feelings_italy : int 67 50 55 50...
> . . .
> $ discrimination_roma : int 7 5 3 3 1 7...
> $ china_Covid_conspiracy : num 2.5 4.25...
> $ generic_Covid_conspiracy: num 2.62 3.25...
```

The Bayesian bootstrapped Kendall's τ correlation matrix for this data is trivially obtained via the `bbcor` function:

```
bb_tau <- bbcor(srol2021, method = "kendall",
  iter = 1000, cores = 1)
bb_summary <- summary(bb_tau,
  ci = 0.9, decimals = 2)
```

Here, the `bbcor` function samples the posterior for the correlation matrix, and takes as arguments the data, the desired correlation type, the number of samples to draw, and the number of cores to use when parallel computing is employed. Printing the returned object outputs the mean correlation matrix. Running `summary` on the returned object and specifying the desired credible interval returns a data frame summarising the posterior with means, standard deviations, and bounds for the credible intervals. For instance, previewing the summary object with `head(bb_summary)` prints the results seen in Listing 1.

²The `BBcor` package can be downloaded from CRAN or from <https://github.com/donaldRwilliams/BBcor>.

**Listing 1 ■ Results of head(bb_summary)**

```
>
> 1      neg_feelings_china--social_distance_china      0.16    0.03    0.10    0.21
> 2      neg_feelings_china--discrimination_china      0.19    0.03    0.13    0.24
> 3      social_distance_china--discrimination_china    0.15    0.04    0.09    0.20
> 4      neg_feelings_china--neg_feelings_italy        0.43    0.03    0.38    0.49
> 5      social_distance_china--neg_feelings_italy      0.06    0.04    0.00    0.11
> 6      discrimination_china--neg_feelings_italy      0.17    0.03    0.11    0.22
```

Depending on the precision of measurements being considered, it can be desirable to obtain more than two decimal points (Cousineau, 2020). This can easily be done by adjusting the decimals setting in the summary method. The posterior means for the correlations and respective intervals can easily be visualized using syntax from the `ggplot2` library (Wickham, 2016). For example, if we subset the data to only include the prejudice and discrimination measures for China and the two conspiracy theory variables, then the following code returns a plot for the ten resulting correlations which can be seen in Figure 2:

```
library(ggplot2)
bb_tau_china <- bbcor(Covid_china_subset,
  method = "kendall")
plot(bb_tau_china) + theme_bw()
```

Analyzing comparisons

The Bayesian bootstrapped correlations can be compared using the `compare` function. The correlations to be compared can be specified either using a character string or by providing a contrast matrix as detailed in section *Comparing Correlations*. For example, if the focus is on comparing the correlation between negative feelings towards China and belief in China-specific Covid-19 conspiracies to the correlation between negative feelings towards China and belief in generic Covid-19 conspiracies, then one can specify the following,

```
comparison <- "
  neg_feelings_china--china_Covid_conspiracy >
  neg_feelings_china--generic_Covid_conspiracy"
compare(comparison, obj = bb_tau_china)
```

which yields a summary of the comparison when printed.

```
> Call:
> lin_comb.bbcor(lin_comb = lin_comb, obj = obj,
  ci = ci, rope = rope,
>   contrast = contrast)
> -----
> Combinations:
> C1: neg_feelings_china--china_Covid_conspiracy
>   > neg_feelings_china--generic_Covid_conspiracy
> -----
> Posterior Summary:
>
> Post.mean Post.sd Cred.lb Cred.ub Prless Prgreat
```

```
> C1 0.12 0.03 0.07 0.17 0 1
> -----
> Note:
> Prless: Posterior probability less than zero
> Prgreat: Posterior probability greater than zero
```

Above, the `comparison` object is a string that states the comparison to be made is that `neg_feelings_china-china_Covid_conspiracy` is greater than `neg_feelings_china-generic_Covid_conspiracy`. This string is passed along to the `compare` function along with the name of the object containing the correlations. The output displays several summary statistics for the posterior of this comparison such as the mean difference, standard deviation, credible interval bounds, and the proportion of posterior mass that is greater or less than zero. In this case, the difference between the two correlations is 0.12, 90% CrI [0.07, 0.17], and the entirety of the posterior mass is above zero.

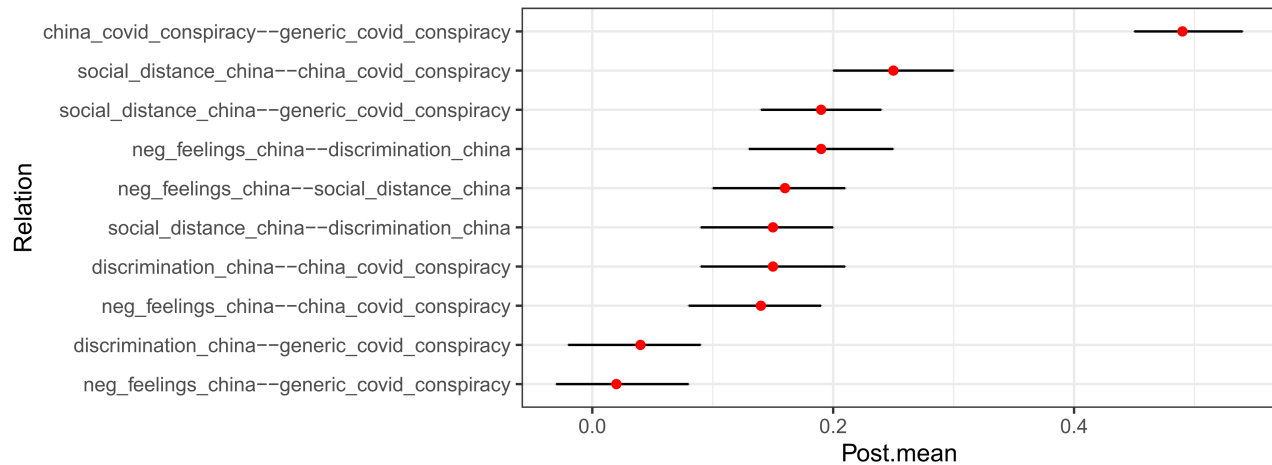
Often, analysts are interested in making more than one comparison. For example, Šrol et al. (2021) repeated the same comparison as above for each country (China, Roma, and Italy) and for each measure of prejudice and discrimination. Thus, there were three comparisons made per country. To avoid tediously typing long character strings, it can be useful to specify a contrast matrix to encode the comparisons of interest. For the subset of variables for China, we must specify a 3×10 matrix corresponding to the three comparisons and ten unique correlations. Additionally, a region of practical equivalence (ROPE) may be stipulated as above, say [-0.10, 0.10]. In R, the analogous code is written as follows

```
contrast_vec <- c(0, 0, 0, 1, 0, 0, -1, 0, 0, 0,
  0, 0, 0, 0, 1, 0, 0, -1, 0, 0,
  0, 0, 0, 0, 0, 1, 0, 0, -1, 0)

contrast_mat <- matrix(contrast_vec, nrow = 3,
  ncol = 10, byrow = TRUE)

compare(obj = bb_tau_china, contrast = contrast_mat,
>   ci = 0.9, rope = c(-0.10, 0.10))
> -----
> Call:
> lin_comb.bbcor(lin_comb = lin_comb, obj = obj,
>   ci = ci, rope = rope, contrast = contrast)
> -----
> Combinations:
```

Figure 2 ■ Output of the `plot` method for Bayesian bootstrapped (BB) correlations obtained with the `bbcor` function. The red dots indicate BB posterior means for the correlations and the bars denote their respective BB 90% credible intervals.



```
> C1: C1
> C2: C2
> C3: C3
> -----
> Posterior Summary:
>
> ROPE: [ -0.1 , 0.1 ]
>
>   Post.mean Post.sd Cred.lb Cred.ub Pr.in
> C1      0.12    0.03   0.07   0.17 0.2762
> C2      0.06    0.03   0.02   0.11 0.9162
> C3      0.11    0.03   0.06   0.16 0.3544
> -----
> Note:
> Pr.in: Posterior probability in ROPE
```

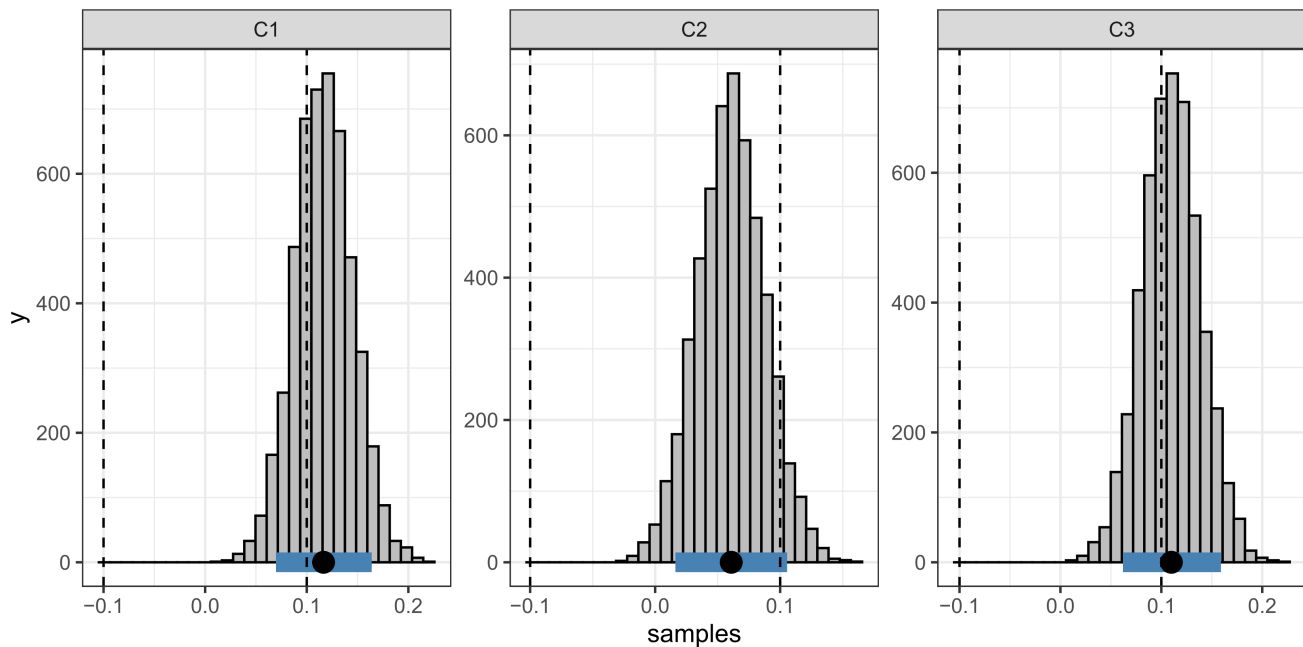
In this output, there are three rows of summary statistics, corresponding to the three comparisons specified in `contrast_mat`. The column `Pr.in` contains the proportion of the posterior mass contained in the ROPE. These combinations can also be visualized via a plotting method. If the `compare` object is saved into an object named `china_comparison`, then calling `plot(china_comparison)` produces the plot in Figure 3. As can be seen, the intervals for all combinations overlap with the ROPE and thus the evidence is ambiguous as to whether these correlations differ. This is a slightly different conclusion than the original analysis where the null hypothesis of no difference was rejected for all three combinations at an $\alpha = 0.05$ level. The results for these comparisons, along with the ones for Italy and Roma can be seen in Table 1.

Recall that the original analysis computed Pearson's correlations, for which there are many tests to probe the difference between two correlations (Diedenhofen &

Musch, 2015). The data, however, were measured using ordinal scales. As such, a measure of association like Kendall's τ may be more useful than a linear association, but this introduces a separate problem due to the lack of a standard error for the difference between two Kendall's τ 's. In our example above, the BB methodology was used to overcome this issue as we trivially estimated and compared the Kendall's τ estimates.

Numerically, the computed correlations and their comparisons were similar to the original, but the resulting interpretations differed. With respect to the magnitude of the associations, some correlations were roughly the same as their Pearson's counterparts, but the majority were slightly weaker, with the differences between the Pearson's and Kendall's τ values ranging from 0.01 to 0.08. Despite these discrepancies, the estimates for the magnitude of the differences were approximately equal between the Pearson and Kendall's correlations. On the other hand, the interpretation of results between the z -test and the ROPE approach diverged. For example, the original analysis failed to reject the null hypothesis for all three comparisons involving Italy. Again, this does not allow statements to be made in support of equality between both correlations. In contrast, the 90% intervals for all three comparisons were trapped completely between $[-0.1, 0.1]$ and under this decision rule, we can conclude that there is evidence to support the respective correlations as practically equivalent. The credible intervals for the remaining comparisons all overlapped with the ROPE and thus there is no decisive evidence for or against equality of the correlations. This also differed from the original analysis in that the majority of

Figure 3 ■ Output from plotting comparisons with the `compare` function. The histograms represent BB posterior samples for the comparisons, black dots indicate the BB mean, and blue bars denote BB 90% CrIs. The dotted black lines capture the bounds for the ROPE of $[-0.1, 0.1]$.



these tests were rejected.

Discussion

In this article we aimed to show how the Bayesian bootstrap can be applied to obtain Bayesian posteriors for correlation coefficients. We began with a concise introduction to the Bayesian bootstrap and provided formulations to obtain Bayesian bootstrapped versions of the Pearson, Spearman's, Gaussian rank, Kendall's, and polychoric correlation coefficients. The main advantage of the BB is that it considerably simplifies obtaining the posterior for the full correlation matrix. A method for comparing correlations was then introduced based on the region of practical equivalence (ROPE) approach (Kruschke & Liddell, 2018). In an empirical application, we demonstrate how a typical analysis of correlations may be carried out using the `BBcor` package. We supplied R code to 1) estimate and visualize posterior estimates for the correlations discussed in the paper and 2) compare correlations using the ROPE approach and visualize the posterior for their difference. Consequently, this example also serves as a tutorial for readers who wish to implement the methodology outlined in this article.

The methods we proposed in this paper contribute to two bodies of literature. The majority of work in psychol-

ogy examining correlations within a Bayesian framework has focused on hypothesis testing with the Bayes factor and thus attention is typically restricted to estimating one or two correlations at a time instead of the full correlation matrix. A complementary view in psychology has called for an increased focus on parameter estimation (Kruschke & Liddell, 2018; Rouder, Haaf, & Vandekerckhove, 2018). Introducing the BB for correlations adds to the literature on Bayesian inference of correlations with a focus on parameter estimation because it is a flexible method capable of estimating the full correlation matrix for a variety of correlation types, and can easily be extended beyond those examined in this paper. Second, a considerable amount of work has been done examining methods for comparing correlations, but this work is focused almost exclusively on the Pearson correlation. By providing a framework wherein a variety of correlations may be compared, the present article also adds to this literature. This is especially valuable for correlations involving ordinal data because of their ubiquity in the social-behavioral sciences.

In our view, the BB is a valuable tool that is best applied when the goal at hand is explicitly parameter estimation as opposed to Bayesian hypothesis testing. Much work at the intersection of psychology and Bayesian statistics has held an eye towards the Bayes factor (BF) for the latter

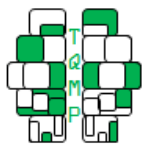


Table 1 ■ Bayesian bootstrapped posterior mean estimates and 90% credible intervals (brackets) for Kendall's τ correlations between measures of discrimination and prejudice and belief in conspiracy theories.

	China-specific covid-19 CTs	Generic covid-19 CTs	Mean Difference
Negative Feelings (China)	0.14 [0.09, 0.19]	0.02 [−0.03, 0.07]	0.12 [0.70, 0.17]
Social Distance (China)	0.25 [0.19, 0.30]	0.19 [0.14, 0.24]	0.06 [0.02, 0.11]
Discrimination (China)	0.15 [0.19, 0.20]	0.04 [−0.02, 0.09]	0.11 [0.06, 0.16]
Negative Feelings (Italy)	0.07 [0.02, 0.12]	0.06 [0.00, 0.11]	0.01 [−0.03, 0.06]
Social Distance (Italy)	0.22 [0.17, 0.28]	0.19 [0.14, 0.24]	0.04 [−0.01, 0.08]
Discrimination (Italy)	0.15 [0.10, 0.21]	0.13 [0.07, 0.18]	0.02 [−0.03, 0.07]
Negative Feelings (Roma)	0.06 [0.01, 0.12]	0.14 [0.08, 0.19]	−0.07 [−0.12, −0.03]
Social Distance (Roma)	0.14 [0.08, 0.19]	0.14 [0.08, 0.19]	0.00 [−0.05, 0.04]
Discrimination (Roma)	0.16 [0.10, 0.22]	0.22 [0.16, 0.27]	−0.06 [−0.11, −0.01]

purpose. Although BFs can be derived using the BB posterior (Newton & Raftery, 1994), it is suboptimal because it depends upon the harmonic mean; a method long known to be problematic (Diciccio, Kass, Raftery, & Wasserman, 1997; Lenk, 2009). If one wants to use the BB to make a decision with respect to a null parameter value, then we view the ROPE approach as a reasonable way of doing so. Moreover, one may want to use an alternative, informed prior when testing a hypothesis. This is challenging with the Bayesian bootstrap because many hyperparameters must be introduced in order to accomplish this (e.g., Poirier, 2011). Thus, the BB shines in the exploratory stages of research because it employs an uninformative prior and can be used to quickly estimate the posterior for the full correlation matrix.

It is important to keep in mind certain aspects of the BB in practice. Practitioners should be wary of applying the BB to small samples (e.g., $n = 10$), as the resulting credible intervals may be more narrow than those obtained, say, with an MCMC Bayesian analysis. The width of the intervals are crucial when using the ROPE approach as overly narrow intervals may result in overconfident inferences. This may be due to the questionable assumption pointed out by Rubin (1981) that values for unobserved data receive no prior, and hence, no posterior, support (but see Hjort, 1991). Thus, if a sample does not include observations from the tails of the population under study (as might often be the case in small samples), then the variance of the BB posterior may be underestimated. However, as n increases, credible intervals based on the BB posterior will converge on those obtained using traditional Bayesian techniques (assuming a uniform prior).

Further the BB diverges from traditional Bayesian methods in some important ways. Most notably, no prior is explicitly elicited by the analyst. Although the subjective choice of specifying a prior can be seen as a core component of Bayesian inference (Savage, 1954), it is often desirable to eliminate this subjectivity in prior spec-

ification (Berger, 2006; Ghosh, 2011). It is interesting to note that Bayesian methods are often favored because they are more consistent with the likelihood principle (Berger & Wolpert, 1988): all the evidence in a sample that is relevant to model parameters is contained in the likelihood function. But the BB violates this principle because the estimation of parameters relies on aggregating datasets which were not observed. An advantageous difference of the BB lies in the computational speed. Many common methods for Bayesian inference are based on MCMC sampling. Because these draws are serially dependent, many samples are typically required for a consistent estimate of the posterior. On the other hand, samples drawn using the BB are independent and thus fewer of them are required. Despite these differences, the BB is a reliable procedure for obtaining a valid posterior distribution.

Conclusion

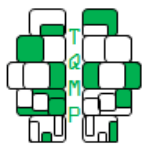
We discussed a generic and simple approach to obtaining posterior distributions via the Bayesian bootstrap (BB) for a variety of correlation coefficients. It is generic because it can be applied broadly to different measures of associations and simple because it amounts to calculating weighted correlations. We further discussed a flexible approach to comparing correlations, or linear combinations thereof. Altogether, the BB provides a powerful tool for approximate Bayesian inference of popular correlation types in the social-behavior sciences.

Authors' note

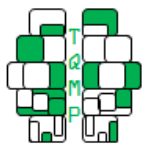
DRW was supported by a National Science Foundation Graduate Research Fellowship under Grant No. 1650042.

References

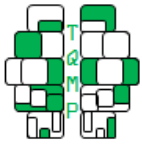
- Albert, J. H. (1992). Bayesian estimation of the polychoric correlation coefficient. *Journal of Statistical Computation and Simulation*, 44(1-2), 47–61. doi:10.1080/00949659208811448



- Bailey, P., Emad, A., Zhang, T., Xie, Q., & Sikali, E. (2018). Weighted and unweighted correlation methods for large-scale educational assessment: wCorr formulas. AIR-NAEP working paper no. 2018-01. NCES data R project series #02. *American Institutes for Research*.
- Barnard, J., McCulloch, R., & Meng, X.-L. (2000). Modeling Covariance Matrices in Terms of Standard Deviations and Correlations, With Application to Shrinkage. *Statistica Sinica*, 10(4), 1281–1311.
- Berger, J. O. (2006). The case for objective Bayesian analysis. *Bayesian Analysis*, 1(3), 385–402. doi:[10.1214/06-BA115](https://doi.org/10.1214/06-BA115)
- Berger, J. O., & Wolpert, R. L. (1988). *The Likelihood Principle*. IMS.
- Boudt, K., Cornelissen, J., & Croux, C. (2012). The Gaussian rank correlation estimator: Robustness properties. *Statistics and Computing*, 22(2), 471–483. doi:[10.1007/s11222-011-9237-0](https://doi.org/10.1007/s11222-011-9237-0)
- Carpenter, B., Gelman, A., Hoffman, M. D., Lee, D., Goodrich, B., Betancourt, M., ... Riddell, A. (2017). Stan: A Probabilistic Programming Language. *Journal of Statistical Software*, 76(1), 1–32. doi:[10.18637/jss.v076.i01](https://doi.org/10.18637/jss.v076.i01)
- Chang, W., Cheng, J., Allaire, J., Sievert, C., Schloerke, B., Xie, Y., ... Borges, B. (2021). *Shiny: Web application framework for R*. Manual.
- Chen, P. Y., Smithson, M., & Popovich, P. M. (2002). *Correlation: Parametric and nonparametric measures*. Sage.
- Chib, S., & Greenberg, E. (1998). Analysis of multivariate probit models. *Biometrika*, 85(2), 347–361. doi:[10.1093/biomet/85.2.347](https://doi.org/10.1093/biomet/85.2.347)
- Choudhuri, N. (1998). Bayesian bootstrap credible sets for multidimensional mean functional. *Annals of statistics*, 2104–2127. doi:[10.1214/aos/1024691463](https://doi.org/10.1214/aos/1024691463)
- Cohen, J., Cohen, P., West, S. G., & Aiken, L. S. (2013). *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*. Routledge.
- Cousineau, D. (2020). How many decimals? rounding descriptive and inferential statistics based on measurement precision. *Journal of Mathematical Psychology*, 97, 102362. doi:[10.1016/j.jmp.2020.102362](https://doi.org/10.1016/j.jmp.2020.102362)
- Devroye, L. (1986). *Non-Uniform Random Variate Generation*. Springer Science & Business Media.
- Diciccio, T. J., Kass, R. E., Raftery, A., & Wasserman, L. (1997). Computing Bayes Factors by Combining Simulation and Asymptotic Approximations. *Journal of the American Statistical Association*, 92(439), 903–915. doi:[10.1080/01621459.1997.10474045](https://doi.org/10.1080/01621459.1997.10474045)
- Diedenhofen, B., & Musch, J. (2015). Cocor: A Comprehensive Solution for the Statistical Comparison of Correlations. *PLoS ONE*, 10(4). doi:[10.1371/journal.pone.0121945](https://doi.org/10.1371/journal.pone.0121945)
- Dunn, O. J., & Clark, V. (1969). Correlation Coefficients Measured on the Same Individuals. *Journal of the American Statistical Association*, 64(325), 366–377. doi:[10.1080/01621459.1969.10500981](https://doi.org/10.1080/01621459.1969.10500981)
- Efron, B. (1979). Bootstrap Methods: Another Look at the Jackknife. *Annals of Statistics*, 7(1), 1–26. doi:[10.1214/aos/1176344552](https://doi.org/10.1214/aos/1176344552)
- Gasparini, M. (1995). Exact multivariate Bayesian bootstrap distributions of moments. *The Annals of Statistics*, 23(3), 762–768. doi:[10.1214/aos/1176324620](https://doi.org/10.1214/aos/1176324620)
- Gelman, A. (2006). Prior distributions for variance parameters in hierarchical models (comment on article by Browne and Draper). *Bayesian Analysis*, 1(3), 515–534. doi:[10.1214/06-BA117A](https://doi.org/10.1214/06-BA117A)
- Ghosh, M. (2011). Objective Priors: An Introduction for Frequentists. *Statistical Science*, 26(2), 187–202. doi:[10.1214/10-STS338](https://doi.org/10.1214/10-STS338)
- Haldane, J. B. S. (1932). A note on inverse probability. *Mathematical Proceedings of the Cambridge Philosophical Society*, 28(1), 55–61. doi:[10.1017/S0305004100010495](https://doi.org/10.1017/S0305004100010495)
- Hjort, N. L. (1991). Bayesian and empirical Bayesian bootstrapping. *Preprint series. Statistical Research Report* <http://urn.nb.no/URN:NBN:no-23420>.
- Jöreskog, K. G. (1994). On the estimation of polychoric correlations and their asymptotic covariance matrix. *Psychometrika*, 59(3), 381–389. doi:[10.1007/BF02296131](https://doi.org/10.1007/BF02296131)
- Kendall, M. G. (1938). A New Measure of Rank Correlation. *Biometrika*, 30(1/2), 81–93. doi:[10.2307/2332226](https://doi.org/10.2307/2332226)
- Kim, Y., & Lee, J. (2003). Bayesian bootstrap for proportional hazards models. *Annals of Statistics*, 31(6), 1905–1922. doi:[10.1214/aos/1074290331](https://doi.org/10.1214/aos/1074290331)
- Krishnamoorthy, K., & Xia, Y. (2007). Inferences on correlation coefficients: One-sample, independent and correlated cases. *Journal of Statistical Planning and Inference*, 137(7), 2362–2379. doi:[10.1016/j.jspi.2006.08.002](https://doi.org/10.1016/j.jspi.2006.08.002)
- Kruschke, J. K. (2018). Rejecting or Accepting Parameter Values in Bayesian Estimation. *Advances in Methods and Practices in Psychological Science*, 1(2), 270–280. doi:[10.1177/2515245918771304](https://doi.org/10.1177/2515245918771304)
- Kruschke, J. K., Aguinis, H., & Joo, H. (2012). The Time Has Come: Bayesian Methods for Data Analysis in the Organizational Sciences. *Organizational Research Methods*, 15(4), 722–752. doi:[10.1177/1094428112457829](https://doi.org/10.1177/1094428112457829)
- Kruschke, J. K., & Liddell, T. M. (2018). The Bayesian New Statistics: Hypothesis testing, estimation, meta-analysis, and power analysis from a Bayesian perspective. *Psychonomic Bulletin & Review*, 25(1), 178–206. doi:[10.3758/s13423-016-1221-4](https://doi.org/10.3758/s13423-016-1221-4)
- Lawrence, E., Bingham, D., Liu, C., & Nair, V. N. (2008). Bayesian Inference for Multivariate Ordinal Data Us-



- ing Parameter Expansion. *Technometrics*, 50(2), 182–191. doi:[10.1198/004017008000000064](https://doi.org/10.1198/004017008000000064)
- Lazar, N. A. (2003). Bayesian empirical likelihood. *Biometrika*, 90(2), 319–326. doi:[10.1093/biomet/90.2.319](https://doi.org/10.1093/biomet/90.2.319)
- Lenk, P. (2009). Simulation Pseudo-Bias Correction to the Harmonic Mean Estimator of Integrated Likelihoods. *Journal of Computational and Graphical Statistics*, 18(4), 941–960. doi:[10.1198/jcgs.2009.08022](https://doi.org/10.1198/jcgs.2009.08022)
- Lewandowski, D., Kurowicka, D., & Joe, H. (2009). Generating random correlation matrices based on vines and extended onion method. *Journal of Multivariate Analysis*, 100(9), 1989–2001. doi:[10.1016/j.jmva.2009.04.008](https://doi.org/10.1016/j.jmva.2009.04.008)
- Lo, A. Y. (1987). A Large Sample Study of the Bayesian Bootstrap. *Annals of Statistics*, 15(1), 360–375. doi:[10.1214/aos/1176350271](https://doi.org/10.1214/aos/1176350271)
- Lo, A. Y. (1988). A Bayesian Bootstrap for a Finite Population. *The Annals of Statistics*, 16(4), 1684–1695. doi:[10.1214/aos/1176351061](https://doi.org/10.1214/aos/1176351061)
- Lyddon, S. P., Holmes, C. C., & Walker, S. G. (2019). General Bayesian updating and the loss-likelihood bootstrap. *Biometrika*, 106(2), 465–478. doi:[10.1093/biomet/asz006](https://doi.org/10.1093/biomet/asz006)
- Meng, X.-L., Rosenthal, R., & Rubin, D. B. (1992). Comparing correlated correlation coefficients. *Psychological Bulletin*, 111(1), 172–175. doi:[10.1037/0033-2909.111.1.172](https://doi.org/10.1037/0033-2909.111.1.172)
- Mulder, J. (2016). Bayes factors for testing order-constrained hypotheses on correlations. *Journal of Mathematical Psychology*. Bayes Factors for Testing Hypotheses in Psychological Research: Practical Relevance and New Developments, 72, 104–115. doi:[10.1016/j.jmp.2014.09.004](https://doi.org/10.1016/j.jmp.2014.09.004)
- Mulder, J., & Pericchi, L. R. (2018). The Matrix- Φ Prior for Estimating and Testing Covariance Matrices. *Bayesian Analysis*, 13(4), 1193–1214. doi:[10.1214/17-BA1092](https://doi.org/10.1214/17-BA1092)
- Newton, M. A., & Raftery, A. E. (1994). Approximate Bayesian Inference with the Weighted Likelihood Bootstrap. *Journal of the Royal Statistical Society. Series B (Methodological)*, 56(1), 3–48. doi:[10.1111/j.2517-6161.1994.tb01956.x](https://doi.org/10.1111/j.2517-6161.1994.tb01956.x)
- Olsson, U. (1979). Maximum likelihood estimation of the polychoric correlation coefficient. *Psychometrika*, 44(4), 443–460. doi:[10.1007/BF02296207](https://doi.org/10.1007/BF02296207)
- Owen, A. (1990). Empirical likelihood ratio confidence regions. *The Annals of Statistics*, 18(1), 90–120. doi:[10.1214/aos/1176347494](https://doi.org/10.1214/aos/1176347494)
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling. *Working Papers*, 8.
- Poirier, D. J. (2011). Bayesian Interpretations of Heteroskedastic Consistent Covariance Estimators Using the Informed Bayesian Bootstrap. *Econometric Reviews*, 30(4), 457–468. doi:[10.1080/07474938.2011.553542](https://doi.org/10.1080/07474938.2011.553542)
- Pozzi, F., Di Matteo, T., & Aste, T. (2012). Exponential smoothing weighted correlations. *The European Physical Journal B*, 85(6), 175. doi:[10.1140/epjb/e2012-20697-x](https://doi.org/10.1140/epjb/e2012-20697-x)
- R Core Team. (2021). *R: A language and environment for statistical computing*. Manual. R Foundation for Statistical Computing. Vienna, Austria.
- Raghunathan, T. E., Rosenthal, R., & Rubin, D. B. (1996). Comparing correlated but nonoverlapping correlations. *Psychological Methods*, 1(2), 178–183. doi:[10.1037/1082-989X.1.2.178](https://doi.org/10.1037/1082-989X.1.2.178)
- Rouder, J. N., Haaf, J. M., & Vandekerckhove, J. (2018). Bayesian inference for psychology, part IV: Parameter estimation and Bayes factors. *Psychonomic Bulletin & Review*, 25(1), 102–113. doi:[10.3758/s13423-017-1420-7](https://doi.org/10.3758/s13423-017-1420-7)
- Rubin, D. B. (1981). The Bayesian Bootstrap. *Annals of Statistics*, 9(1), 130–134. doi:[10.1214/aos/1176345338](https://doi.org/10.1214/aos/1176345338)
- Savage, L. J. (1954). *The foundations of statistics*. John Wiley & Sons (Oxford, England).
- Šrol, J., Cavojava, V., & Mikušková, E. B. (2021). Social consequences of COVID-19 conspiracy beliefs: Evidence from two studies in Slovakia. PsyArXiv. doi:[10.31234/osf.io/y4svc](https://doi.org/10.31234/osf.io/y4svc)
- Steiger, J. H. (1980). Tests for comparing elements of a correlation matrix. *Psychological Bulletin*, 87(2), 245–251. doi:[10.1037/0033-2909.87.2.245](https://doi.org/10.1037/0033-2909.87.2.245)
- Talhok, A., Doucet, A., & Murphy, K. (2012). Efficient Bayesian Inference for Multivariate Probit Models With Sparse Inverse Correlation Matrices. *Journal of Computational and Graphical Statistics*, 21(3), 739–757. doi:[10.1080/10618600.2012.679239](https://doi.org/10.1080/10618600.2012.679239)
- Tokuda, T., Goodrich, B., Van Mechelen, I., Gelman, A., & Tuerlinckx, F. (2011). Visualizing distributions of covariance matrices. *Columbia Univ., New York, USA, Tech. Rep*, 18–18.
- van Doorn, J., Ly, A., Marsman, M., & Wagenmakers, E.-J. (2018). Bayesian Inference for Kendall's Rank Correlation Coefficient. *The American Statistician*, 72(4), 303–308. doi:[10.1080/00031305.2016.1264998](https://doi.org/10.1080/00031305.2016.1264998)
- Vandekerckhove, J., Rouder, J. N., & Kruschke, J. K. (2018). Editorial: Bayesian methods for advancing psychological science. *Psychonomic Bulletin & Review*, 25(1), 1–4. doi:[10.3758/s13423-018-1443-8](https://doi.org/10.3758/s13423-018-1443-8)
- Wagenmakers, E.-J., Marsman, M., Jamil, T., Ly, A., Verhaegen, J., Love, J., ... Morey, R. D. (2018). Bayesian inference for psychology. Part I: Theoretical advantages



- and practical ramifications. *Psychonomic Bulletin & Review*, 25(1), 35–57. doi:[10.3758/s13423-017-1343-3](https://doi.org/10.3758/s13423-017-1343-3)
- Wagenmakers, E.-J., Morey, R. D., & Lee, M. D. (2016). Bayesian Benefits for the Pragmatic Researcher. *Current Directions in Psychological Science*, 25(3), 169–176. doi:[10.1177/0963721416643289](https://doi.org/10.1177/0963721416643289)
- Wagenmakers, E.-J., Verhagen, J., & Ly, A. (2016). How to quantify the evidence for the absence of a correlation. *Behavior Research Methods*, 48(2), 413–426. doi:[10.3758/s13428-015-0593-0](https://doi.org/10.3758/s13428-015-0593-0)
- Webb, E. L., & Forster, J. J. (2008). Bayesian model determination for multivariate ordinal and binary data. *Computational Statistics & Data Analysis*, 52(5), 2632–2649. doi:[10.1016/j.csda.2007.09.008](https://doi.org/10.1016/j.csda.2007.09.008)
- Weng, C.-S. (1989). On a Second-Order Asymptotic Property of the Bayesian Bootstrap Mean. *The Annals of Statistics*, 17(2), 705–710. doi:[10.1214/aos/1176347136](https://doi.org/10.1214/aos/1176347136)
- Wetzels, R., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for correlations and partial correlations. *Psychonomic Bulletin & Review*, 19(6), 1057–1064. doi:[10.3758/s13423-012-0295-x](https://doi.org/10.3758/s13423-012-0295-x)
- Wickham, H. (2016). *Ggplot2: Elegant graphics for data analysis*. Springer-Verlag New York.
- Williams, D. R., & Mulder, J. (2019). *Bayesian Hypothesis Testing for Gaussian Graphical Models: Conditional Independence and Order Constraints*. PsyArXiv. doi:[10.31234/osf.io/ypxd8](https://doi.org/10.31234/osf.io/ypxd8)
- Yuan, Y., & Johnson, V. E. (2008). Bayesian Hypothesis Tests Using Nonparametric Statistics. *Statistica Sinica*, 18(3), 1185–1200.
- Zou, G. Y. (2007). Toward using confidence intervals to compare correlations. *Psychological Methods*, 12(4), 399–413. doi:[10.1037/1082-989X.12.4.399](https://doi.org/10.1037/1082-989X.12.4.399)

Appendix

Following Rubin (1981), let $\mathbf{d} = (d_1, \dots, d_K)'$ be the vector of all K possible distinct values in $\mathbf{x} = (x_1, \dots, x_n)'$ and let $\boldsymbol{\theta} = (\theta_1, \dots, \theta_K)'$ be a vector of probabilities associated with \mathbf{d} such that

$$p(x_i = d_k | \boldsymbol{\theta}) = \theta_k, \quad i = 1 \dots, n; \quad k = 1, \dots, K, \quad (24)$$

and the sum of all probabilities equal one. If \mathbf{x} is an i.i.d. sample from (24) and n_k is the number of values in \mathbf{x} equal to d_k , then the prior for $\boldsymbol{\theta}$ under the Bayesian bootstrap is the so called Haldane prior (Haldane, 1932)

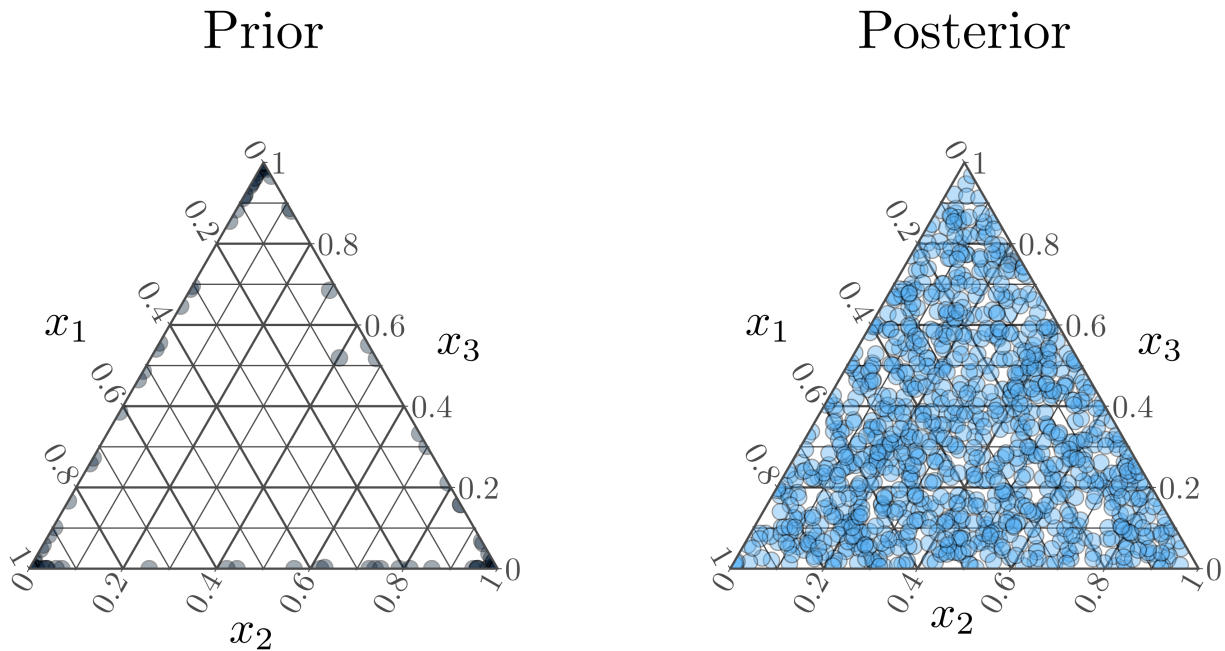
$$p(\boldsymbol{\theta}) \propto \prod_{k=1}^K \theta_k^{-1}, \quad (25)$$

and corresponds to the improper prior Dirichlet distribution $\text{Dir}(\boldsymbol{\alpha})$ with $\boldsymbol{\alpha} = (0, \dots, 0)$. When this prior is combined with a multinomial likelihood, it yields a posterior for $\boldsymbol{\theta}$ which follows the Dirichlet distribution with $\boldsymbol{\alpha} = (1, \dots, 1)$, that is,

$$\begin{aligned} p(\boldsymbol{\theta} | \mathbf{x}) &\propto p(\mathbf{x} | \boldsymbol{\theta}) p(\boldsymbol{\theta}) \\ &\propto \prod_{k=1}^K \theta_k^{n_k} \prod_{k=1}^K \theta_k^{-1} \\ &\propto \prod_{k=1}^K \theta_k^{n_k - 1}. \end{aligned} \quad (26)$$

A BB prior distribution (using $\alpha_i = 0.1$) and a corresponding posterior distribution are plotted in Figure 4. As can be seen, the prior mass is mostly placed over probabilities near zero and one. In the limit, as all $\alpha_i \rightarrow 0$, there is zero mass placed over θ 's for unobserved data. The posterior distribution places mass uniformly on $[0, 1]$ which indicates that any combination of θ 's for the observed values is equally likely.

Figure 4 ■ Ternary plots of the prior (left) and posterior (right) distributions for the parameter θ under the Bayesian bootstrap for three observations.



Citation

Rodriguez, J. E., & Williams, D. R. (2022). An example of the use of the TQMP class. *The Quantitative Methods for Psychology*, 18(1), 39–54. doi:[10.20982/tqmp.18.1.p039](https://doi.org/10.20982/tqmp.18.1.p039)

Copyright © 2022, *Rodriguez and Williams*. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) or licensor are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

Received: 28/10/2021 ~ Accepted: 20/12/2021